Alibaba Cloud ApsaraDB for MySQL

User Guide

Issue: 20190415

MORE THAN JUST CLOUD | 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 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

Legal disclaimer..... I
Generic conventions..... I
1 Preface.....1
2 Quick start..... 3
3 Data migration..... 5
3.1 Migrate from other cloud databases to ApsaraDB RDS.....5
3.2 Use mysqldump to migrate MySQL data.....5
3.3 Migrate RDS data to the local database.....9
3.3.1 Migrate RDS for MySQL data to a local MySQL database..... 9
3.3.2 Migrate RDS for SQL Server data to a local SQL Server database.....12
3.3.3 Migrate RDS for PostgreSQL data to a local PostgreSQL database.... 15
3.3.4 Migrate RDS for PPAS data to a local Oracle database..... 16
3.3.5 Migrate RDS for PPAS data to a local PPAS database..... 20
3.4 Compress data..... 21
3.5 Use psql to migrate PostgreSQL data.....22
3.6 Migrate SQL Server to cloud.....24
3.6.1 Migrate data to RDS for SQL Server 2008 R2.....24
3.6.2 Migrate data to RDS for SQL Server 2012/2016.....31
3.7 Migrate a MySQL database from Tencent Cloud to Alibaba Cloud..... 41
3.8 Migrate a MySQL database from Google Cloud to Alibaba Cloud.....47
4 Billing management.....55
4.1 Change the billing method..... 55
4.2 Enable auto-renewal for a Subscription instance.....56
4.3 Manually renew a Subscription instance.....57
5 Instance management.....59
5.1 Restart an instance..... 59
5.2 Configure the maintenance period..... 59
5.3 Migrate instance across zones..... 60
5.4 Switch between master and slave instances..... 66
5.5 Modify the data replication mode.....68
5.6 Create a read-only instance..... 70
5.7 Release an instance.....72
5.8 Upgrade the database version.....73
5.9 RDS for MySQL release notes.....74
5.10 Change configurations.....76
5.11 SQL Server DBCC function.....79
5.12 End connections for SQL Server instances.....80
6 Read/write splitting..... 81
6.1 Introduction to read/write splitting..... 81

6.2 Enable read/write splitting....84
6.3 Modify the latency threshold and read weights....88
6.4 Switch read/write splitting address type....90
6.5 Disable read/write splitting....91
6.6 Monitor read/write splitting performance....92
6.7 Test read/write splitting performance....93
6.8 Verify read weight distribution....97
6.9 Rules of weight distribution by system....98
7 Account management....101
7.1 Reset the instance password....101
7.2 Change account permissions....101
7.3 Authorize a service account....102
7.4 Delete an account....104
7.5 Manage the LOGIN user for SQL Server instances....105
7.6 Manage users for SQL Server instances....106
8 Database management....108
8.1 Create a database....108
8.2 Manage databases of SQL Server instances....108
9 Connection management....111
9.1 Set the access mode....111
9.2 Set network type....112
9.3 Hybrid access solution for smooth migration from classic networks to VPCs....114
9.4 Set intranet and Internet IP addresses....119
10 Monitoring and Alarming....125
10.1 Set the monitoring frequency....125
10.2 Set monitoring rules....130
11 Security....133
11.1 SQL audit....133
11.2 Switch the IP whitelist to enhanced security mode....135
11.3 Set the whitelist....136
11.4 Set SSL encryption....139
11.5 Set TDE....143
12 Log management....145
13 SQL explorer....147
14 Backup....152
14.1 Back up RDS data....152
14.2 View the free quota of the backup space....157
14.3 Download data and log backup files....158
14.4 Logical backup and recovery for PPAS....160
15 Recovery....162
15.1 Restore MySQL data....162

15.2 Restore MySQL databases or tables..... 171
15.3 Restore SQL Server Data..... 174
15.4 Restore PostgreSQL or PPAS data..... 182
15.5 Restore MariaDB data..... 191
15.6 Recover data to a temporary instance (RDS for SQL Server)..... 200
16 Typical applications..... 204
16.1 Cached data persistence..... 204
16.2 Multi-structure data storage..... 205

1 Preface

Overview

ApsaraDB for Relational Database Service (RDS) is a stable and reliable online database service with auto-scaling capabilities. Based on Apsara distributed file system and high-performance SDD storage, RDS supports MySQL, SQL Server, PostgreSQL, and PPAS engines, and provides a complete set of solutions for disaster recovery, backup, recovery, monitoring, migration, and others. This helps you operate and manage your own database. For benefits of RDS, see *Benefits*.

This document describes RDS features and functions and further explains the procedure to configure RDS through the *RDS console* . You can also manage RDS through APIs and SDKs.

If you need technical assistance, you can open the *RDS console* and choose Support > Open a new ticket or *click here* to submit a ticket.

For more information about functions and pricing of RDS, log on to *official website of ApsaraDB for RDS*.

Declaration

Some features or services described in this document may be unavailable for certain regions. See the relevant commercial contracts for specific terms and conditions.

This document serves as a user guide. No content in this document can constitute any express or implied warranty.

The content of this document is updated based on product upgrade and many other factors. You must first verify the document with your latest software version.

Consideration

RDS supports multiple types of databases. This document takes MySQL as an example to describe the features and usage of RDS. Some types of databases may not support certain features. The actual interface may vary slightly.

General terms

- Instance: A database service process that takes up physical memory independently . You can set different memory size, disk space, and database type, among which

the memory specification determines the performance of the instance. After the instance is created, you can change the configuration and delete the instance at any time.

- Database: A logical unit created in an instance. Multiple databases can be created in an instance, and the database name is unique within the instance.
- Region and zone: A region is a physical data center. A zone is a physical area that has independent power supply and networks within a region. For more information, see *Alibaba Cloud Global Infrastructure*.

Common conventions

Term	Description
Local database/Source database	Refers to the database deployed in the local equipment room or the database not on the ApsaraDB. In most cases, it refers to the source database to be migrated to the ApsaraDB in this document.
RDS for XX (MySQL, SQL Server, PostgreSQL, PPAS)	It indicates the RDS of a specific database type, for example, RDS for MySQL means the instance enabled on the RDS with a database type of MySQL.

2 Quick start

If you use RDS for the first time, see the following *Quick Start* documents to get started with RDS.

- *Quick Start for MySQL*
- *Quick Start for SQL Server*
- *Quick Start for PostgreSQL*
- *Quick Start for PPAS*

If you have questions beyond *Quick Start*, see *User Guide*.

Database engines

ApsaraDB for MySQL

MySQL is the world's most popular open source database. As an important part of LAMP and a combination of open source software (Linux + Apache + MySQL + Perl/PHP/Python), MySQL is widely used in a variety of applications.

In the Web 2.0 era, MySQL serves as the basis of the underlying architecture of the popular BBS software system Discuz! and blogging platform WordPress. In the Web 3.0 era, leading Internet companies including Alibaba, Facebook, and Google have built their large-scale mature database clusters by taking advantage of the advanced flexibility of MySQL.

Based on Alibaba's MySQL source code branch, ApsaraDB for MySQL proves to have excellent performance and throughput. It withstands the massive data traffic and a large number of concurrent users during many November 11 (Singles' Day) shopping festivals - the Chinese equivalent of Cyber Monday. ApsaraDB for MySQL also offers a range of advanced functions including optimized read/write splitting, data compression, and intelligent optimization.

RDS for MySQL currently supports versions 5.5, 5.6, and 5.7.

ApsaraDB for SQL Server

SQL Server is one of the first commercial databases and is an important part of the Windows platform (IIS + .NET + SQL Server), with support for a wide range of enterprise applications. The SQL Server Management Studio software comes with a

rich set of built-in graphical tools and script editors. You can quickly get started with a variety of database operations through visual interfaces.

Powered by a high-availability architecture and the capability to recover data at any point in time, ApsaraDB for SQL Server provides strong support for a variety of enterprise applications. It also covers Microsoft’s licensing fee.

RDS for SQL Server currently supports the following versions:

- SQL Server 2008 R2 Enterprise
- SQL Server 2012 Web, Standard, and Enterprise
- SQL Server 2016 Web, Standard, and Enterprise

ApsaraDB for PostgreSQL

PostgreSQL is the world’s most advanced open source database. As an academic relational database management system, it provides full compliance with SQL specifications and robust support for a diverse range of data formats (including JSON , IP, and geometric data, which are not supported by most commercial databases).

ApsaraDB for PostgreSQL supports a range of features including transactions, subqueries, Multi-Version Concurrency Control (MVCC), and data integrity verificati on. It also integrates a number of important functions, including high availability, backup, and recovery, to help mitigate your O&M burden.

RDS for PostgreSQL currently supports version 9.4.

ApsaraDB for PPAS

Postgres Plus Advanced Server (PPAS) is a stable, secure, and scalable enterprise-level relational database. Based on PostgreSQL, PPAS delivers enhanced performance, application solutions, and compatibility, and provides the capability to run Oracle applications directly. It is a reliable and cost-effective option for running a variety of enterprise applications.

ApsaraDB for PPAS incorporates a number of advanced functions including account management, resource monitoring, backup, recovery, and security control, and it continues to be updated and improved regularly.

RDS for PPAS currently supports version 9.3.

3 Data migration

3.1 Migrate from other cloud databases to ApsaraDB RDS

You can smoothly migrate data from other cloud databases to ApsaraDB RDS.

Migrate MySQL from AWS RDS to ApsaraDB RDS

Migrate MySQL from AWS RDS to ApsaraDB RDS with DTS

3.2 Use mysqldump to migrate MySQL data

mysqldump is used to migrate MySQL data. The disadvantage of mysqldump is that the service downtime is long. Use mysqldump if the data volume is small or if a long service downtime is allowed.

Background information

As RDS is fully compatible with MySQL, the procedure for migrating local databases to an RDS instance is similar to the procedure for migrating data from one MySQL server to another.

Prerequisites

- You have set a whitelist, applied for an Internet IP address, and created databases and accounts for the RDS instance. For more information, see *Quick Start*.
- An ECS instance has been created.

Procedure

Before data migration, create a migration account in the local database, and grant read and write permissions of the database to the migration account.

1. Create a migration account in the local database.

```
CREATE USER ' username '@' host ' IDENTIFIED BY ' password
';
```

Parameter description:

- `username` : indicates the account to be created.
- `host` : indicates the host from which you log on to the database using the account. As a local user, you can use `localhost` to log on to the database. To log on from any host, you can use the wildcard `%`.
- `password` : indicates the logon password of the account.

In the following example, the user name is `William` and password is `Changme123`. The user is allowed to log on to the local database from any host.

```
CREATE USER ' William '@'%' IDENTIFIED BY ' Changme123 ';
```

2. Grant permissions to the migration account in the local database.

```
GRANT SELECT ON databasena me . tablename TO ' username
'@' host ' WITH GRANT OPTION ;
```

```
GRANT REPLICATIO N SLAVE ON databasena me . tablename
TO ' username '@' host ' WITH GRANT OPTION ;
```

Parameter description:

- `privileges` : indicates the operating authorization of the account, such as *SELECT*, *INSERT*, and *UPDATE*. To grant all permissions to the account, use *ALL*.
- `databasena me` : indicates the database name. To grant all database permissions to the account, use the wildcard *.
- `tablename` : indicates the table name. To grant all table permissions to the account, use the wildcard *.
- `username` : indicates the name of the account to be granted permissions.
- `host` : indicates the host authorized for the account to log on to the database. As a local user, you can use *localhost* to log on to the database. To log on from any host, you can use the wildcard *%*
- `WITH GRANT OPTION` : indicates an optional parameter that enables the account to use the `GRANT` command.

In the following example, the account *William* is granted all database and table permissions:

```
GRANT ALL ON *.* TO ' William '@'%';
```

3. Use the data export tool of mysqldump to export data in the database as data files.

> **Note:**
> Do not update data during data export. This step is used to export data only, excluding stored procedures, triggers, and functions.

```
mysqldump - h localIp - u userName - p -- opt -- default -
character - set = utf8 -- hex - blob dbName -- skip - triggers
> / tmp / dbName . sql
```

Parameter description:

- `localIp` : IP address of the local database server
- `userName` : Migration account of the local database
- `dbName` : Name of the database to be migrated
- `/ tmp / dbName . sql` : Backup file name

4. Use mysqldump to export stored procedures, triggers, and functions.

   Note:
   If no stored procedures, triggers, and functions are used in the database, you may skip this step. When exporting stored procedures, triggers, and functions, you must remove `definer` for compatibility with RDS.

   ```
   mysqldump -h localIp -u userName -p --opt --default
   -character-set=utf8 --hex-blob dbName -R | sed -e
   's/DEFINER[ ]*=[ ]*[^*]*\*/\*/' > /tmp/triggerProcedure.
   sql
   ```

   Parameter description:

   - `localIp` : IP address of the local database server
   - `userName` : Migration account of the local database
   - `dbName` : Name of the database to be migrated
   - `/tmp/triggerProcedure.sql` : Backup file name

5. Upload the data files and stored procedure files to ECS.

   The example in this document describes how to upload files to the following path.

   ```
   /tmp/dbName.sql
   /tmp/triggerProcedure.sql
   ```

6. Log on to ECS and import data files and stored procedure files to the target RDS.

   ```
   mysql -h intranet4example.mysql.rds.aliyuncs.com -u
     userName -p dbName < /tmp/dbName.sql
   mysql -h intranet4example.mysql.rds.aliyuncs.com -u
     userName -p dbName < /tmp/triggerProcedure.sql
   ```

   Parameter description:

   - `intranet4example.mysql.rds.aliyuncs.com` : RDS instance connection address. An intranet IP address is used as an example.
   - `userName` : Migration account of the RDS database
   - `dbName` : Name of the database to be imported
   - `/tmp/dbName.sql` : Name of the data file to be imported
   - `/tmp/triggerProcedure.sql` : Name of the stored procedure file to be imported

3.3 Migrate RDS data to the local database

3.3.1 Migrate RDS for MySQL data to a local MySQL database

RDS for MySQL supports migration of cloud data to local databases using physical and logical backup files.

Export based on a physical backup file

Background information

Due to software restrictions, data recovery is supported only in Linux currently. If you want to recover data to Windows, first you need recover data to Linux and then migrate the data to Windows.

Prerequisites

Data restoration tool Percona XtraBackup has been installed in the Linux system.

- For MySQL 5.6 and earlier version, install Percona XtraBackup 2.3.
- For MySQL 5.7, install Percona XtraBackup 2.4.

For installation instructions, see *Percona XtraBackup 2.3* and *Percona XtraBackup 2.4*.

Procedure

This example assumes that the local server runs the RHEL6/x64 system and the path for saving the backup file is */home/mysql/*.

1. Download the physical backup file and upload the file to the target server. For more information about how to obtain the backup file, see *Download RDS data and log backup*. If the target server can access the source instance, you can use `wget " url "` to download the backup file. *url* indicates the backup file downloading address.
2. Switch to the backup file path.

   ```
   cd / home / mysql /
   ```

3. Decompress the backup file.

   ```
   tar   vizxf   filename . tar . gz
   ```

   `filename.tar.gz` indicates the name of the backup file.

4. Check whether the databases contained in the decompressed file are correct.

   ```
   cd   filename /
   ```

```
ll
```

The system displays the following information, in which *db0dz1rv11f44yg2*, *mysql*, and *test* are databases in RDS:

```
- rw - r -- r --  1   root   root         269   Aug   19   18 : 15
 backup - my . cnf
 drwxr - xr - x   2   root   root        4096   Aug   21   10 : 31
 db0dz1rv11  f44yg2
- rw - rw ----  1   root   root   209715200   Aug    7   10 : 44
 ibdata1
 drwxr - xr - x   2   root   root        4096   Aug   21   10 : 31
 mysql
 drwxr - xr - x   2   root   root        4096   Aug   21   10 : 31
 test
- rw - r -- r --  1   root   root          10   Aug   19   18 : 15
 xtrabackup  _binary
- rw - r -- r --  1   root   root          23   Aug   19   18 : 15
 xtrabackup  _binlog_in  fo
- rw - r -- r --  1   root   root          77   Aug   19   18 : 15
 xtrabackup  _checkpoin  ts
- rw - r -- r --  1   root   root        2560   Aug   19   18 : 15
 xtrabackup  _logfile
- rw - r -- r --  1   root   root          72   Aug   19   18 : 15
 xtrabackup  _slave_inf  o
```

5. Recover the data file.

```
 innobackup  ex  -- defaults - file =./ backup - my . cnf  -- apply
 - log  ./
```

Data is successfully recovered when the system displays `innobackup ex : completed OK !`.

6. Modify the configuration file. In the *backup-my.cnf* file, comment out *innodb_fas t_checksum*, *innodb_page_size*, and *innodb_log_block_size*, and add *datadir=/home/mysql*, as shown in the following example.

```
#  This   MySQL   options   file   was   generated   by
 innobackup  ex - 1 . 5 . 1 .
#  The   MySQL   Server
[ mysqld ]
 innodb_dat  a_file_pat  h = ibdata1 : 200M : autoextend
 innodb_log  _files_in_  group = 2
 innodb_log  _file_size = 524288000
# innodb_fas  t_checksum = 0
# innodb_pag  e_size = 16364
# innodb_log  _block_siz  e = 512
 datadir =/ home / mysql /
```

7. Reinstall MySQL and obtain the root permission of the database.

```
 rm  - rf   mysql
```

```
mysql_inst  all_db  -- user = mysql  -- datadir =/ home / mysql /
```

If the system displays the following information, the mysql system table is successfully reinstalled.

```
Installing  MySQL  system  table ...
OK
Filling  help  table ...
OK
```

8. Modify the file owner.

```
chown  - R  mysql : mysql  / home / mysql /
```

9. Start the mysqld process.

```
mysqld_saf  e  -- defaults - file =/ home / mysql / backup - my .
cnf  &
```

10. Log on to the database from a client.

```
mysql - u  root  - p
```

11. Verify database integrity.

```
show  databases ;
```

The database is successfully recovered when the system displays the following information:

```
+--------------------+
|  Database          |
+--------------------+
|  informatio  n_schema  |
|  db0dz1rv11  f44yg2    |
|  mysql             |
|  performanc  e_schema  |
|  test              |
+--------------------+
```

Export based on a logical backup file

This example assumes that the local server runs the RHEL6/x64 system and the path for saving the backup file is */home/mysql/*

Procedure

1. Download logical backup file and upload the file to the target server. For more information about how to obtain the backup file, see *Download RDS data and log backup*. If the target server can access the source instance, you can use `wegt  "`

`url "` to download the backup file. *url* indicates the backup file downloading address.

2. Switch to the backup file path.

```
cd / home / mysql /
```

3. Decompress the backup file.

```
tar vizxf filename . tar . gz
```

filename.tar.gz indicates the name of the backup file.

4. Decompress the SQL file.

```
gunzip filename . sql . gz
```

filename.sql.gz indicates the name of the compressed SQL file.

5. Perform logical import to import data to the target database.

```
mysql - u userName - p - h hostName - P port dbName <
filename . sql
```

filename.sql indicates the name of the decompressed SQL file.

3.3.2 Migrate RDS for SQL Server data to a local SQL Server database

RDS for SQL Server supports migration of cloud data to local databases using physical backup files.

Procedure

1. Download the full and incremental physical backup files of RDS and upload the files to the target server.

   For more information about how to obtain the backup file, see *Download RDS data and log backup*.

   If the target server can access the source instance, you can use `wegt " url "` to download the backup file. *url* indicates the backup file downloading address.

2. Decompress the full physical backup file and incremental physical backup file.

   A backup file is named in the format of database name+backup type+date and time +task ID.bak, of which backup type may be one of the following:

   - datafull: indicates full backup, such as rdsumu2myfzbeai1_datafull_2014 02250050_2250050.bak.
   - datadiff: indicates incremental backup, such as rdsumu2myfzbeai1_dat adiff_201402260050_2260050.bak.
   - log: indicates log backup, such as rdsumu2myfzbeai1_log_201402260050_226005 0.bak.

3. Obtain the decompressed full backup file and incremental backup file. This example assumes that the backup files are stored in the following paths:

   - Path for saving the full backup file: *d :\ backup \ rdsumu2myf zbeai1_dat afull_2014 02250050_2 250050 . bak*
   - Path for saving the incremental backup file: *d :\ backup \ rdsumu2myf zbeai1_dat adiff_2014 02260050_2 260050 . bak*

4. Log on to the local SQL Server console and query the logical names of the RDS files based on the backup files.

```
restore   fileliston  ly   from   disk =& apos ; d :\ backup \
rdsumu2myf  zbeai1_dat  afull_2014  02250050_2  250050 . bak & apos
;
```

```
go
```

The system displays the following information, where the logical name of the data file is data1 and that of the log file is log.

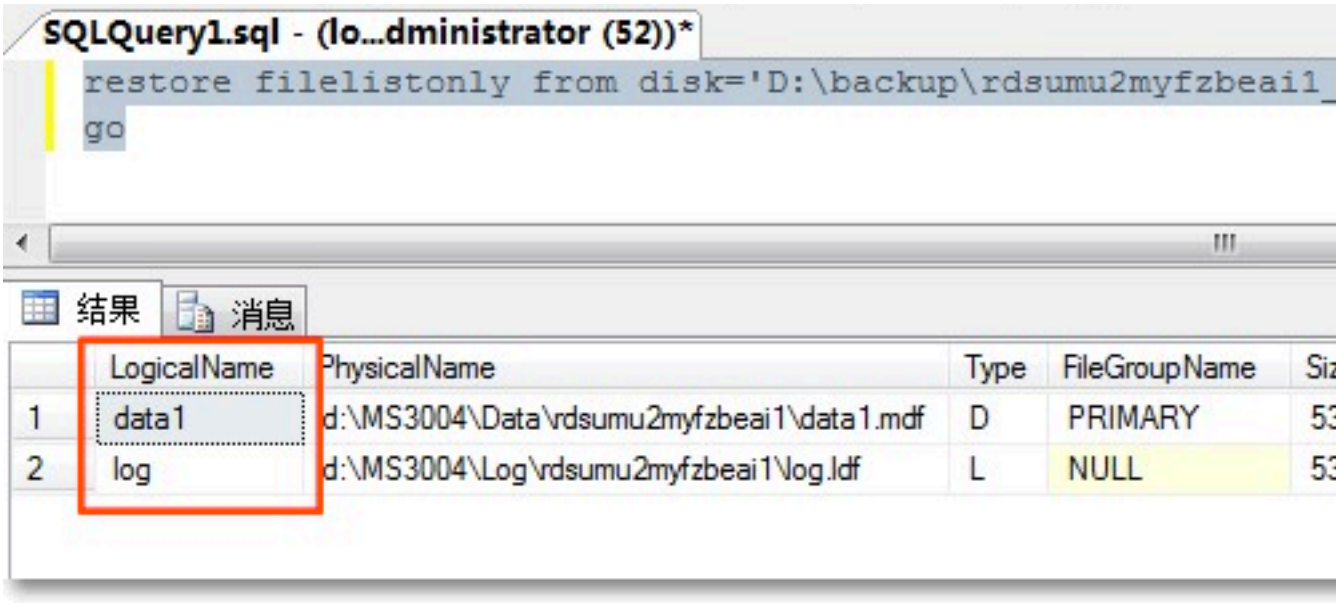

5. Load the full backup file.

```
 restore   database   rdsumu2myf  zbeai1   from   disk =& apos ; d :
\ backup \ rdsumu2myf  zbeai1_dat  afull_2014  02250050_2  250050 .
 bak & apos ;  with     replace , norecovery , stats = 10 ,
 move  & apos ; data1 & apos ;  to  & apos ; d :\ database \
 rdsumu2myf  zbeai1 \ data \ data1 . mdf & apos ;,
 move  & apos ; log & apos ;  to  & apos ; d :\ database \
 rdsumu2myf  zbeai1 \ log \ log . ldf & apos ;
 go
```

Parameters description:

- *d:\database\rdsumu2myfzbeai1\data* is the data address, and *data1.mdf* is the logical name of the data file
- *d:\database\rdsumu2myfzbeai1\log* is the log address, and *log.ldf* is the logical name of the log file

After the script is executed, database *rdsumu2myfzbeai1* is in Recovering state.

Note:
If you only want to recover full backup data, skip Step 6 and proceed to Step 7. If you also want to recover incremental backup data, perform Step 6.

6. Load the incremental backup file.

```
restore database rdsumu2myf zbeai1 from disk =& apos ; D :
\ backup \ rdsumu2myf zbeai1_dat adiff_2014 02260050_2 260050 .
bak & apos ; with replace , norecovery , stats = 10 ,
move & apos ; data1 & apos ; to & apos ; d :\ database \
rdsumu2myf zbeai1 \ data \ data1 . mdf & apos ;,
move & apos ; log & apos ; to & apos ; d :\ database \
rdsumu2myf zbeai1 \ log \ log . ldf & apos ;
go
```

After the script is executed, database *rdsumu2myfzbeai1* is in Recovering state.

7. Recover the database.

```
restore database rdsumu2myf zbeai1 with recovery
go
```

After the script is executed, database *rdsumu2myfzbeai1* is available.

3.3.3 Migrate RDS for PostgreSQL data to a local PostgreSQL database

RDS for PostgreSQL supports migration of cloud data to local databases using logical backup files.

Procedure

1. Connect the PostgreSQL client to RDS.
2. Run the following command to back up data.

```
pg_dump -U username -h hostname -p port databasena
me -f filename
```

Parameters description:

- `username` : indicates the user name used for database logon.
- `hostname` : indicates the host name of the database.
- `port` : indicates the database port number.
- `databasena me` : indicates the name of the database you want to back up.
- `filename` : indicates the name of the backup file to be generated.

For example:

```
pg_dump -U myuser -h rds2z2tp80 v3752wb455 . pg . rds .
aliyuncs . com -p 3433 pg001 -f pg001 . sql
```

3. Save the *pg001.sql* backup file to the target server.

4. Run the following command to recover data to the local database:

```
psql -U username -h hostname -d desintatio ndb -p
port -f dumpfilena me . sql
```

Parameter description:

- `username` : indicates the user name used for database logon.
- `hostname` : indicates the database address.
- `port` : indicates the database port number.
- `databasena me` : indicates the database name.
- `filename` : indicates the backup file name.

For example:

```
psql -U myuser -h localhost -d pg001 -p 5432 -f
pg001 . sql
```

Since the permission configuration of the RDS database is inconsistent with that of the local database, some permission-related warnings or errors may occur during data import. They can be ignored, for example:

```
WARNING : no privileges could be revoked for " xxxxx "
ERROR : role " xxxxx " does not exist
```

3.3.4 Migrate RDS for PPAS data to a local Oracle database

Constraint

Now only files and normal types of data can be exported. BLOB and other binary types are not supported.

Prerequisites

- An Oracle database must be installed on the server.
- The IP address of the Oracle server must be added to the whitelist of the RDS for PPAS database instance. For specific instructions, see *Set whitelist*.
- You must create a table structure in Oracle that corresponds to the RDS for PPAS database table structure.
- The PostgreSQL client has been uploaded to the Oracle database server.

Procedure

Note:

This document uses the migration of data from RDS for PPAS to an Oracle database installed on an ECS instance as an example. In this example, the ECS instance OS is CentOS 6.5.

1. Install the PostgreSQL client on the Oracle database server.

```
[ root @ oraclexe  ~]#  yum   install   postgresql . x86_64
[ root @ oraclexe  ~]# / usr / bin / psql  -- version
 psql  ( PostgreSQL )  8 . 4 . 20
```

2. On the ECS instance, configure password-free logon for RDS for PPAS.

```
[ root @ oraclexe  ~]#  vim  ~/. pgpass
[ root @ oraclexe  ~]#  cat  ~/. pgpass
 rm - 2ze466l5u1  k657yyn . ppas . rds . aliyuncs . com : 3433 :
ora : myadmin : xxxxxxx
// Parameter   format :  HOSTNAME : PORT : DATABASE : USERNAME :
PASSWORD
[ root @ oraclexe  ~]#  chmod   0600  ~/. pgpass
```

Note:

The configuration file . *pgpass* is located in the *HOME* directory.

3. Test the connection between ECS and RDS for PPAS.

```
[ root @ oraclexe  ~]#  psql  - h   rm - 2ze466l5u1  k657yyn . ppas
. rds . aliyuncs . com  - p   3433  - U   myadmin   ora
 psql . bin  ( 9 . 3 . 1 . 3 ,  server   9 . 3 . 13 . 37 )
 Input  " help "  to   obtain   help   informatio  n .
 ora =>
```

If you can log on to RDS for PPAS as user ora, the connection has been established. After a successful test, return to user root.

```
 ora => \ q
[ root @ oraclexe  ~]#
```

4. Create a data export script in the ECS instance.

   a. Create a file *ppas_exp_a ll_tables_ to_csv . sh* .

```
  vi   ppas_exp_a  ll_tables_  to_csv . sh
```

   b. Insert the following text into the *ppas_exp_a ll_tables_ to_csv . sh* script:

```
#  ppas_exp_a  ll_tables_  to_csv . sh  < hostname > < port > <
 username > < database >
 #  Author :  Xiao   Shaocong  ( Scott   Siu )
 #  Email :  shaocong . xsc @ alibaba - inc . com
  TMP_PATH ="/ tmp / ppas_table  s_ $ 1_ $ 2_ $ 3_ $ 4 "
  mkdir  $ TMP_PATH
  if  [ $? - ne   0  ]
```

```
 then
     exit  1 ;
 fi
 echo  " select  '$ 1  $ 2  $ 3  $ 4  ' ||  tablename  || ' $
TMP_PATH  ' ||  tablename  from  pg_tables  where  tableowner
='$ 3 '  and  ( schemaname ='$ 3 '  or  schemaname =' public
');" > / tmp / ppas_table  s_ $ 1_ $ 2_ $ 3_ $ 4 . sql
 psql  - h  $ 1  - p  $ 2  - U  $ 3  $ 4  - f  / tmp /
ppas_table  s_ $ 1_ $ 2_ $ 3_ $ 4 . sql  |  head  - n  - 2  |
tail  - n  + 3  |  awk  - F  " " '{ printf  (" psql  - h  % s
- p  % s  - U  % s  % s  - c  \"\\ copy  % s  TO  '\''% s /% s
'\''  CSV  HEADER \"\ n ",$ 1 ,$ 2 ,$ 3 ,$ 4 ,$ 5 ,$ 6 ,$ 7 )}'
|  sh
```

5. **Grant the execution permission to the** *ppas_exp_a ll_tables_ to_csv . sh* **script.**

```
[ root @ oraclexe  ~]#  chmod  0755  ppas_exp_a  ll_tables_
to_csv . sh
```

6. **Run the data export script in the ECS instance.**

```
[ root @ oraclexe  ~]# ./ ppas_exp_a  ll_tables_  to_csv . sh
rm - 2ze466l5u1  k657yyn . ppas . rds . aliyuncs . com  3433
myadmin  ora
```

7. **Verify the data in the exported CSV file.**

```
[ root @ oraclexe  ~]#  cat  / tmp / ppas_table  s_rm - 2ze466l5u1
k657yyn . ppas . rds . aliyuncs . com_3433_m  yadmin_ora /*
 deptno , dname , loc
 10 , ACCOUNTING , NEW  YORK
 20 , RESEARCH , DALLAS
 30 , SALES , CHICAGO
 40 , OPERATIONS , BOSTON
 empno , ename , job , mgr , hiredate , sal , comm , deptno
 7369 , SMITH , CLERK , 7902 , 17 - DEC - 80  00 : 00 : 00 , 800 .
00 ,, 20
 7499 , ALLEN , SALESMAN , 7698 , 20 - FEB - 81  00 : 00 : 00 ,
1600 . 00 , 300 . 00 , 30
 7521 , WARD , SALESMAN , 7698 , 22 - FEB - 81  00 : 00 : 00 ,
1250 . 00 , 500 . 00 , 30
 7566 , JONES , MANAGER , 7839 , 02 - APR - 81  00 : 00 : 00 ,
2975 . 00 ,, 20
 7654 , MARTIN , SALESMAN , 7698 , 28 - SEP - 81  00 : 00 : 00 ,
1250 . 00 , 1400 . 00 , 30
 7698 , BLAKE , MANAGER , 7839 , 01 - MAY - 81  00 : 00 : 00 ,
2850 . 00 ,, 30
 7782 , CLARK , MANAGER , 7839 , 09 - JUN - 81  00 : 00 : 00 ,
2450 . 00 ,, 10
 7788 , SCOTT , ANALYST , 7566 , 19 - APR - 87  00 : 00 : 00 ,
3000 . 00 ,, 20
 7839 , KING , PRESIDENT ,, 17 - NOV - 81  00 : 00 : 00 , 5000 .
00 ,, 10
 7844 , TURNER , SALESMAN , 7698 , 08 - SEP - 81  00 : 00 : 00 ,
1500 . 00 , 0 . 00 , 30
 7876 , ADAMS , CLERK , 7788 , 23 - MAY - 87  00 : 00 : 00 , 1100
. 00 ,, 20
 7900 , JAMES , CLERK , 7698 , 03 - DEC - 81  00 : 00 : 00 , 950 .
00 ,, 30
```

```
 7902 , FORD , ANALYST , 7566 , 03 - DEC - 81   00 : 00 : 00 , 3000
. 00 ,, 20
 7934 , MILLER , CLERK , 7782 , 23 - JAN - 82   00 : 00 : 00 , 1300
. 00 ,, 10
 empno , startdate , enddate , job , sal , comm , deptno , chgdesc
 7369 , 17 - DEC - 80   00 : 00 : 00 ,, CLERK , 800 . 00 ,, 20 ,
New   Hire
 7499 , 20 - FEB - 81   00 : 00 : 00 ,, SALESMAN , 1600 . 00 , 300
. 00 , 30 , New   Hire
 7521 , 22 - FEB - 81   00 : 00 : 00 ,, SALESMAN , 1250 . 00 , 500
. 00 , 30 , New   Hire
 7566 , 02 - APR - 81   00 : 00 : 00 ,, MANAGER , 2975 . 00 ,, 20 ,
New   Hire
 7654 , 28 - SEP - 81   00 : 00 : 00 ,, SALESMAN , 1250 . 00 , 1400
. 00 , 30 , New   Hire
 7698 , 01 - MAY - 81   00 : 00 : 00 ,, MANAGER , 2850 . 00 ,, 30 ,
New   Hire
 7782 , 09 - JUN - 81   00 : 00 : 00 ,, MANAGER , 2450 . 00 ,, 10 ,
New   Hire
 7788 , 19 - APR - 87   00 : 00 : 00 , 12 - APR - 88   00 : 00 : 00
, CLERK , 1000 . 00 ,, 20 , New   Hire
 7788 , 13 - APR - 88   00 : 00 : 00 , 04 - MAY - 89   00 : 00 : 00
, CLERK , 1040 . 00 ,, 20 , Raise
 7788 , 05 - MAY - 90   00 : 00 : 00 ,, ANALYST , 3000 . 00 ,, 20 ,
Promoted   to   Analyst
 7839 , 17 - NOV - 81   00 : 00 : 00 ,, PRESIDENT , 5000 . 00 ,, 10
, New   Hire
 7844 , 08 - SEP - 81   00 : 00 : 00 ,, SALESMAN , 1500 . 00 , 0 .
00 , 30 , New   Hire
 7876 , 23 - MAY - 87   00 : 00 : 00 ,, CLERK , 1100 . 00 ,, 20 ,
New   Hire
 7900 , 03 - DEC - 81   00 : 00 : 00 , 14 - JAN - 83   00 : 00 : 00
, CLERK , 950 . 00 ,, 10 , New   Hire
 7900 , 15 - JAN - 83   00 : 00 : 00 ,, CLERK , 950 . 00 ,, 30 ,
Changed   to   Dept   30
 7902 , 03 - DEC - 81   00 : 00 : 00 ,, ANALYST , 3000 . 00 ,, 20 ,
New   Hire
 7934 , 23 - JAN - 82   00 : 00 : 00 ,, CLERK , 1300 . 00 ,, 10 ,
New   Hire
```

8. Import the CSV file into the Oracle database.

   - Method 1: Use Oracle SQL Loader to import data. For more information, see *Oracle SQL Loader Overview*.
   - Method 2: Use Oracle SQL Developer to import data. For more information, see *SQL Developer Concepts and Usage*.

Troubleshooting

Problem

During the execution of data export script, the system displays a message indicating that a directory cannot be created.

```
[ root @ oraclexe  ~]# ./ ppas_exp_a  ll_tables_  to_csv . sh   rm -
 2ze466l5u1  k657yyn . ppas . rds . aliyuncs . com   3433   myadmin
 ora
```

```
mkdir : Cannot create directory : "/ tmp / ppas_table s_rm -
2ze466l5u1 k657yyn . ppas . rds . aliyuncs . com_3433_m yadmin_ora
": file already exists
```

Solution

Delete the existing directory.

```
[ root @ oraclexe ~]# rm - rf / tmp / ppas_table s_rm -
2ze466l5u1 k657yyn . ppas . rds . aliyuncs . com_3433_m yadmin_ora
```

3.3.5 Migrate RDS for PPAS data to a local PPAS database

ApsaraDB for PPAS supports migration of cloud data to local databases using logical backup files.

Procedure

1. **Connect the PostgreSQL client to RDS.**
2. **Run the following command to back up data.**

```
pg_dump - U username - h hostname - p port databasena
me - f filename
```

Parameter descriptions:

- `username` **: indicates the user name used for database logon.**
- `hostname` **: indicates the host name of the database.**
- `port` **: indicates the database port number.**
- `databasena me` **: indicates the name of the database you want to back up.**
- `filename` **: indicates the name of the backup file to be generated. For example:**

```
pg_dump - U ppas_user - h rdsv07z563 m7o25cj550 public .
ppas . rds . aliyuncs . com - p 3433 edb - f ppas . sql
```

3. **Save the** *ppas.sql* **backup file to the target server.**

4. Run the following command to recover data to the local database:

```
psql -U username -h hostname -d desintatio ndb -p
port -f dumpfilena me . sql
```

Parameter descriptions:

- `username` : indicates the user name used for database logon.
- `hostname` : indicates the database address.
- `port` : indicates the database port number.
- `databasena me` : indicates the database name.
- `filename` : indicates the backup file name. For example:

```
psql -U ppas_user -h localhost -d edb -p 5444 -
f ppas . sql
```

As the permission settings of the RDS database are different from those of the local database, some permission-related warnings or errors may occur during data import. They can be ignored, for example:

```
WARNING : no privileges could be revoked for " xxxxx
"
ERROR : role " xxxxx " does not exist
```

3.4 Compress data

RDS for MySQL 5.6 supports data compression through the TokuDB storage engine. A large number of tests showed that, after data tables are switched from the InnoDB storage engine to the TokuDB storage engine, the amount of data can be reduced by 80% to 90%, that is, 2 TB of data can be compressed to 400 GB or even less. The TokuDB storage engine supports transactions and online DDL operations, which are compatible with applications running on a MyISAM or an InnoDB storage engine.

Restrictions

- The TokuDB storage engine does not support foreign keys.
- The TokuDB storage engine is not applicable to scenarios where frequent and massive data read operations are required.

Procedure

1. Run the following command to check the MySQL version.

```
SELECT   version ();
```

> Note:
> Currently, only MySQL 5.6 supports the TokuDB storage engine. As for MySQL 5.1 or 5.5, you have to upgrade it to MySQL 5.6 first.

2. Set the proportion of `loose_toku db_buffer_ pool_ratio` , namely, the proportion that TokuDB occupies in the shared cache of TokuDB and InnoDB.

```
select   sum ( data_lengt  h )  into  @ all_size   from
informatio  n_schema . tables   where   engine =' innodb ';
select   sum ( data_lengt  h )  into  @ change_siz  e   from
informatio  n_schema . tables   where   engine =' innodb '  and
concat ( table_sche  ma , '.',  table_name )  in  (' XX . XXXX ', '
XX . XXXX ', ' XX . XXXX ');
select   round (@ change_siz  e /@ all_size * 100 );
```

   In the preceding code, *XX.XXXX* refers to the database and table to be transferred to the TokuDB storage engine.

3. Restart the instance.

   For more information, see *Restart an instance*.

4. Modify the storage engine.

```
ALTER   TABLE   XX . XXXX   ENGINE = TokuDB
```

   In the preceding code, *XX.XXXX* refers to the database and table to be transferred to the TokuDB storage engine.

3.5 Use psql to migrate PostgreSQL data

This document describes how to use psql commands to restore the PostgreSQL data backup file to the target RDS.

Background information

PostgreSQL supports logical backup. To import PostgreSQL data, use the pg_dump logical backup function to export backup files and then import the files to the RDS through psql.

Prerequisite

You have set a whitelist, applied for an Internet IP address, and created databases and accounts for the RDS instance. For more information, see *Quick Start*.

Prepare local data

1. Connect to the local PostgreSQL database through the PostgreSQL client.
2. Run the following command to back up data:

```
pg_dump -U username -h hostname -p port databasena
me -f filename
```

Parameters are described as follows:

- `username` : User name for the local database
- `hostname` : Local database host name. `localhost` can be used if you log on to the local database host.
- `port` : Local database port number
- `databasena me` : Name of the local database to be backed up
- `filename` : Name of the backup file to be generated

For example, to use the database account William to back up the local PostgreSQL database, log on to the PostgreSQL host and run the following command:

```
pg_dump -U William -h localhost -p 3433 pg001 -f
pg001.sql
```

Migrate data

Note:

Network stability and data security are improved when data is restored through the intranet. We recommend that you upload the data to ECS and then restore the data to the target RDS through the intranet. If a data file is too large, compress it before uploading. This scenario is explained in the following example:

1. Log on to ECS.

2. Run the following command through the PostgreSQL client to import data into the RDS:

```
psql -U username -h hostname -d desintatio ndb -p
port -f dumpfilena me . sql
```

Parameters are described as follows:

- `username` : PostgreSQL database user name on the RDS
- `hostname` : PostgreSQL database address on the RDS
- `port` : PostgreSQL database port number on the RDS
- `databasena me` : PostgreSQL database name on the RDS
- `filename` : Local backup data file name

For example:

```
psql -U William -h postgresql . rds . aliyuncs . com -d
pg001 -p 3433 -f pg001 . sql
```

Since the permission configuration of the RDS database is inconsistent with that of the local database, some permission-related warnings or errors may occur during data import. They can be ignored, for example:

```
WARNING : no privileges could be revoked for " xxxxx "
ERROR : role " xxxxx " does not exist
```

3.6 Migrate SQL Server to cloud

3.6.1 Migrate data to RDS for SQL Server 2008 R2

Instances of the SQL Server 2008 R2 version support easy data migration to the cloud database. You only have to back up all the data using the official backup function of Microsoft on the self-built database, upload the backup file to the *Object Storage Service (OSS)* of Alibaba Cloud, and then move the full amount of data to the specified RDS database through the RDS console. This feature takes advantage of Microsoft's official backup and recovery program, realizes 100% compatibility, and is combined with the powerful capabilities of OSS. All these functions make it a highly efficient feature for data migration to the cloud database.

Prerequisite

A target database has been created in RDS. For more information, see *Create database and account for SQL Server 2008 R2*.

> **Note:**
> The name of the target database in RDS can be the same with that of the local database to be migrated.

Billing details

When you migrate data to the cloud, no additional fees are charged for RDS but you must pay for OSS, as shown in the following figure.

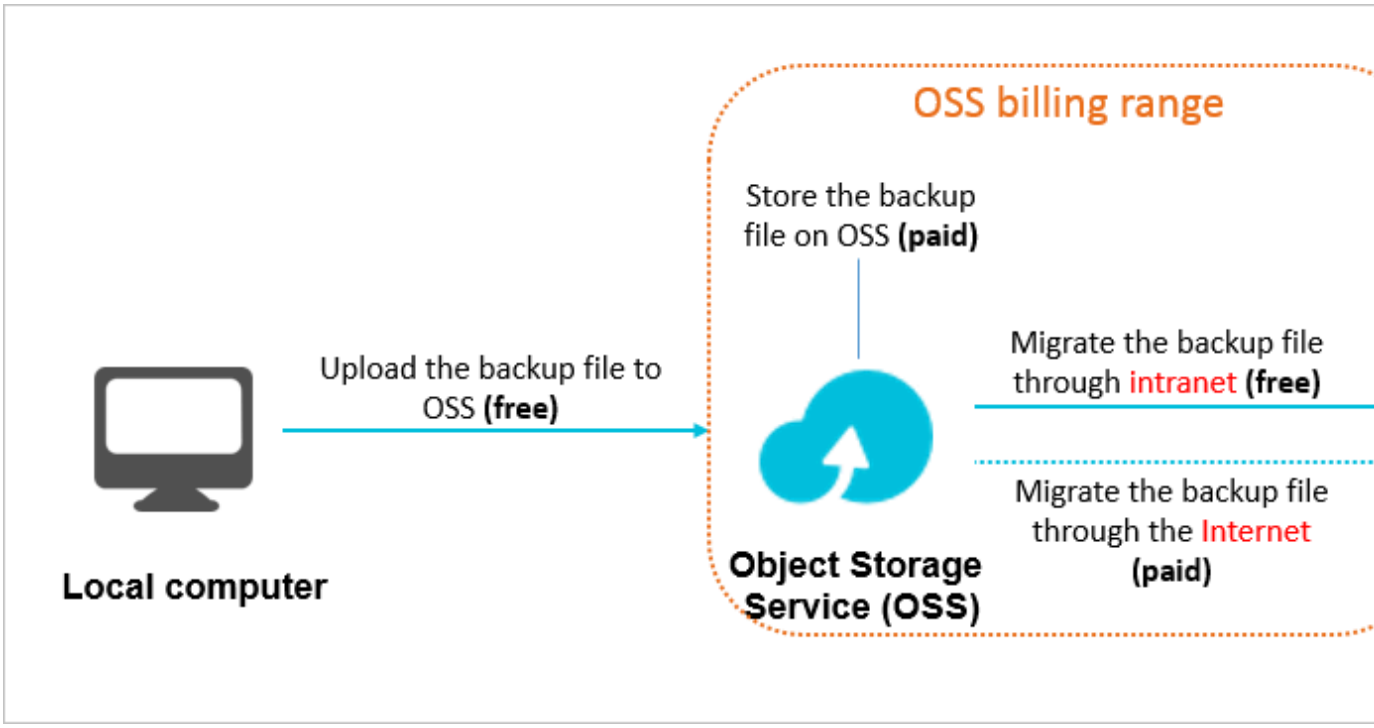

Figure description:

- Uploading local data backup files to OSS is free of charge.
- OSS storage can be changed if you store backup files on OSS. For more information, see *Pricing*.
- If you migrate backup files from OSS to RDS through intranet, no extra fees are charged. If it is through Internet, OSS charges for the Internet outbound traffic. For more information, see *Pricing*.

> **Note:**

The RDS instance and OSS bucket can connect to each other through intranet only when they are located in the same region. Therefore, make sure that the backup files are uploaded to the bucket that is located in the same region as the target RDS instance.

Procedure

1. Prepare the local database. The detailed procedure is as follows:
   a. Start the Microsoft SQL Server Management Studio (SSMS) client.
   b. Log on to the database to be migrated.
   c. Run the following commands to check the recover mode of the local database:

```
use   master ;
go
select   name ,  case   recovery_m  odel
when   1   then   FULL
when   2   then   BULD_LOGGE  D
when   3   then   SIMPLE   end   model   from   sys . databases
where   name   not   in  ( master , tempdb , model , msdb );
go
```

   Check the model value of the local database:

   - If the model value is not `FULL` , go to Step d.
   - If the model value is `FULL` , go to Step e.

   d. Run the following commands to set the recover mode of the source database to `FULL` .

   Note:
   Setting recover mode to `FULL` increases the number of SQL Server logs. Therefore, make sure there is sufficient disk space for the logs.

```
ALTER   DATABASE  [ dbname ]  SET   RECOVERY   FULL ;
go
ALTER   DATABASE  [ dbname ]  SET   AUTO_CLOSE   OFF ;
go
```

   e. Run the following commands to back up the source database. This example uses *filename . bak* as the backup file name.

```
use   master ;
go
BACKUP   DATABASE  [ testdbdb ]  to   disk  = d :\ backup \
filename . bak   WITH   COMPRESSIO  N , INIT ;
```

```
go
```

f. Run the following commands to verify integrity of the backup file.

```
USE   master
GO
RESTORE   FILELISTON  LY
  FROM   DISK  =  ND :\ Backup \ filename . bak ;
```

Returned result description:

- If a result set is returned, the backup file is valid.
- If an error is returned, the backup file is invalid. In this case, back up the database again.

g. Run the following commands to recover the recover mode of the source database.

Note:

If you do not perform Step iv (that is, the original recover mode of the database is `FULL`), skip this step.

```
ALTER   DATABASE  [ dbname ]  SET   RECOVERY   SIMPLE ;
```

```
go
```

2. Upload the local backup file to OSS and retrieve the file URL. The detailed procedure is as follows:

   a. Upload the backup file to OSS:

      - For the procedure of uploading a file smaller than 5 GB, see *Upload an object*.
      - For the procedure of uploading multiple files or a file larger than 5 GB, see *Multipart upload*. To perform this step on GUIs, see *ossbrowser*.

   b. In the left-side navigation pane of the *OSS console*, select the bucket where the backup file belongs.

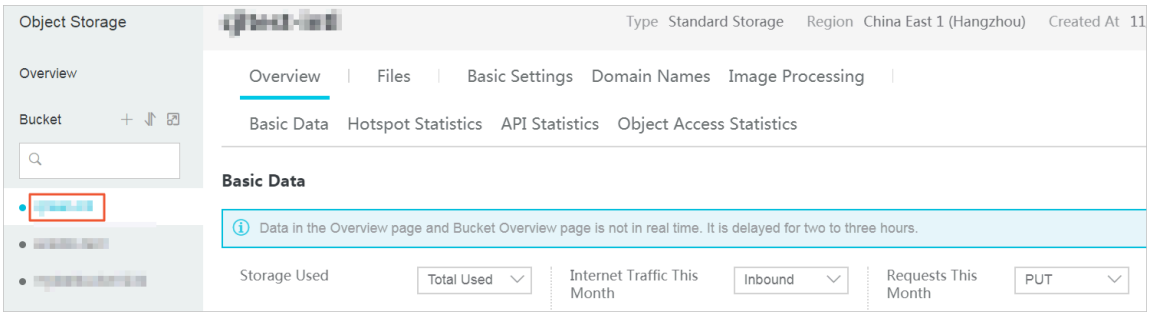

   c. Select Files.

   d. Click the name of the target backup file.

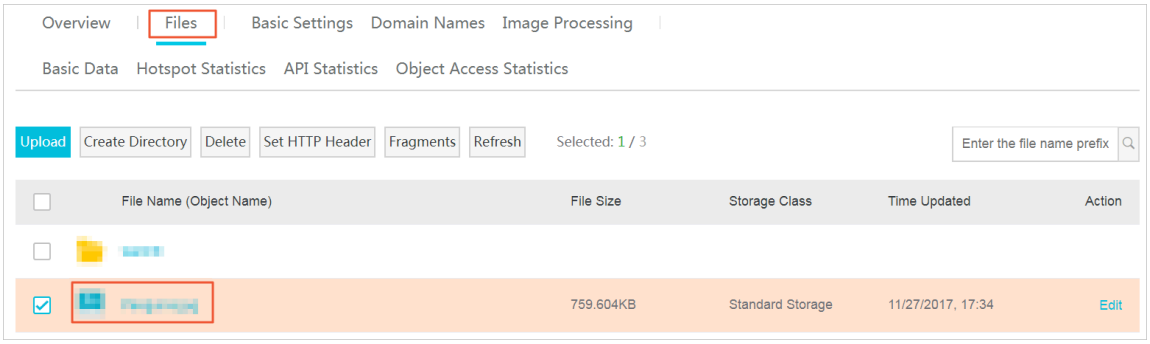

   e. In the Signature field, change the validity period of the link. We recommend that you set the validity period to 28,800s, namely, eight hours.

> **Note:**
> When you migrate the backup file from OSS to RDS, the URL of the backup file is required. If the link validity period for the URL expires, the data

migration fails. Therefore, we recommend that you set the validity period to the maximum value, which is 28,800s.

f. Click Copy File URL. The default URL is the Internet connection address of the file.

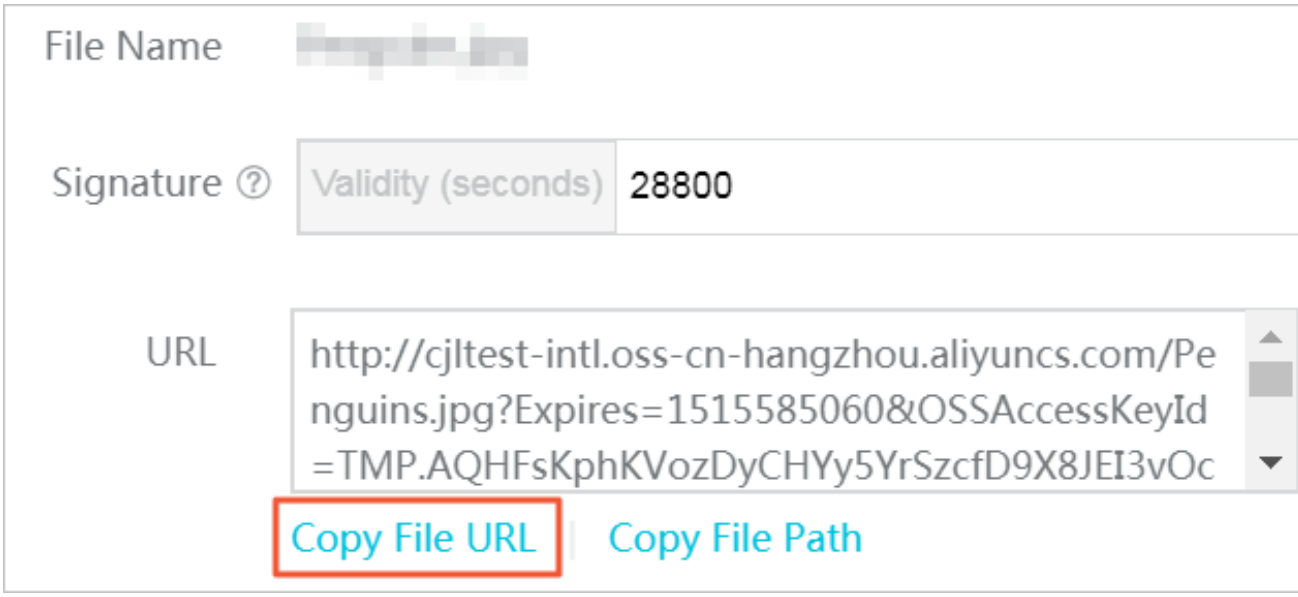

g. If you want to migrate data through the intranet, change the endpoint in the backup file URL to the intranet endpoint. The intranet endpoint varies with the network type and region. For more information, see *Access domain name and data center*.

For example, if the backup file URL is `http://rdstest-yanhua.oss-cn-shanghai.aliyuncs.com/testmigrat erds_20170906143807_FULL.bak?Expires=1514189963&OSSAccessKeyId=TMP.AQGVf994YTPfArSpw78uix2rdGBi-dPe_FzQSLwOLP7MVlR-XXXX`, change the Internet endpoint `oss-cn-shanghai.aliyuncs.com` in the URL to the intranet endpoint `oss-cn-shanghai-internal.aliyuncs.com`.

3. Migrate the backup file from OSS to RDS. The detailed procedure is as follows:
   a. Log on to the *RDS console*.
   b. Select the region where the target instance is located.
   c. Click the ID of the target instance to go to the Basic Information page.
   d. In the left-side navigation pane, click Databases to go to the Databases page.
   e. Find the target database and click Migrate backup files from OSS in the Action column.

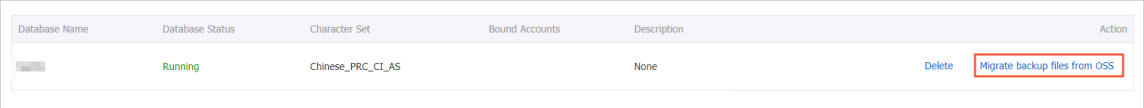

   f. In the Import Guide dialog box, read the prompt and click Next to go to the Upload the backup files page.
   g. Read the prompt and click Next to go to the Import data page.
   h. In the Backup file OSS URL box, enter the backup file URL in OSS.

> **Note:**

Currently, RDS supports only one cloud migration solution, that is one-time migration of the full backup file.

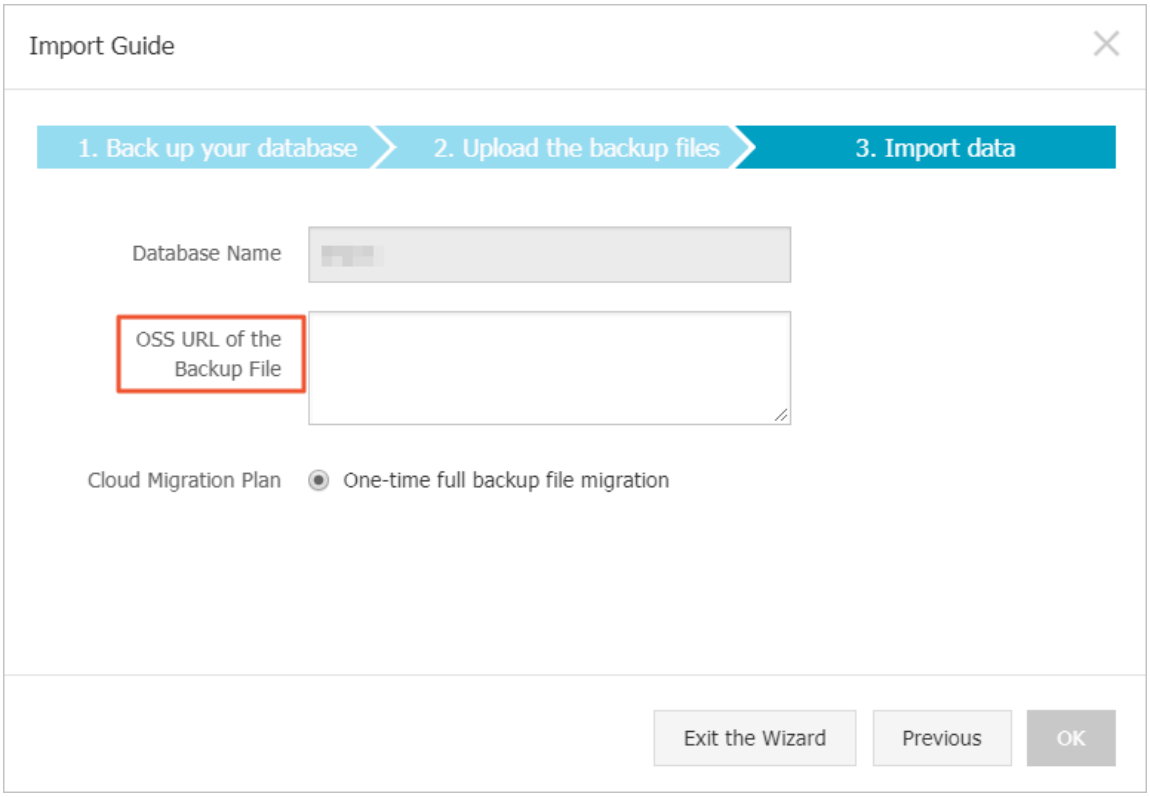

i. Click OK.

j. In the left-side navigation pane, click Data Migration to Cloud to go to the page listing the tasks of migrating backup files from OSS to RDS.

k. Find the target migration task. If Tasks Status is Success, the data is successfully migrated to the RDS database. If the migration task status does not change to Success after a long time, click View File Details next to the migration task to view the failure causes. After solving the problems, perform the required steps to migrate the backup file again.

3.6.2 Migrate data to RDS for SQL Server 2012/2016

This document describes how to migrate full backup data to RDS for SQL Server 2012/2016.

Applicable versions

- Basic series (single-node): RDS for SQL Server 2016/2012 Web or Enterprise Edition
- High-availability series (dual-node): RDS for SQL Server 2016/2012 Standard or Enterprise Edition

For instructions on how to migrate data to RDS for SQL Server 2008 R2 Enterprise Edition (high-availability series), see *Migrate data to RDS for SQL Server 2008 R2*.

Restrictions

Backup file version

Backup data of new SQL Server versions cannot be migrated to old SQL Server versions. For example, you cannot migrate data from SQL Server 2016 to SQL Server 2012.

Backup file type

Differential and log backup files are not supported.

Backup file suffix

The backup file suffix must be bak, diff, trn, or log. If your backup file is not generated using the script provided in this document, use one of the following suffix:

- bak: indicates a full backup file.
- diff: indicates a differential backup file.
- trn or log: indicates a transaction log backup file.

Backup file name

The name of the full backup file cannot contain certain special characters, such as @ or |. Otherwise, the migration will fail.

Precautions

AliyunRDSImportRole

After you authorize the RDS official service account to access OSS, the system creates the role AliyunRDSImportRole in the RAM system. Do not modify or delete the role. Otherwise, the backup upload cannot succeed, and you need to perform the authorization on the wizard again.

Delete backup file from OSS

Before the backup restoration is complete, do not delete the backup file from OSS.

Prerequisites

Instance capacity

Ensure that the RDS for SQL Server instance has sufficient storage space. Expand the space if needed.

A database with the same name is not allowed in the target instance.

You do not need to create a target database in advance. This is different from the requirement stated in *Migrate data to RDS for SQL Server 2008 R2*.

If the target RDS instance already has a database whose name is the same as that of a database to be migrated, back up and delete the database in the target RDS instance before creating a migration task.

Create a superuser account on target instance.

It is recommended that you create a superuser account for the target instance on the console in advance. If the target instance does not have a superuser account, the migration can succeed but you cannot access the database unless you take measures by referring to Common Errors at the end of this document.

For information about how to create a superuser account, see *Create accounts and databases (SQL Server 2012 or 2016)*.

Prepare an OSS bucket

Create an OSS bucket that is in the same region as the target instance if you do not have one.

1. Log on to *OSS console*.
2. Click the + sign on the left pane.

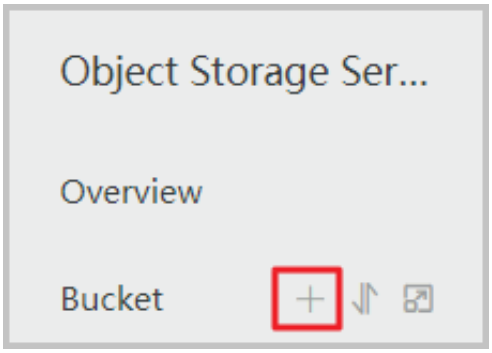

3. Set the bucket name, region, storage class, and ACL permission, and click OK. (Ensure that the bucket is in the same region as the target RDS for SQL Server instance so that the bucket can be selected in subsequent steps.)

Run DBCC CHECKDB

Run DBCC CHECKDB(‘xxx’) on the local database and ensure that the result is as follows, with no allocation errors or consistency errors.

```
CHECKDB found 0 allocation errors and 0 consistenc y
errors in database ' xxx '.
DBCC execution completed . If DBCC printed error
messages , contact your system administra tor .
```

If DBCC CHECKDB shows any errors, fix them before migration.

Procedure

Only three steps are required to migrate a local database to an RDS for SQL Server 2012/2016 instance:

1. Back up the local database.
2. Upload the backup file to OSS.
3. Create a migration task.

Back up the local database

Before performing a full backup of the local database, stop writing data into the database. Data written into the database during the backup will not be backed up.

You can perform a full backup by using your own method or following these steps:

1. Download the *backup script* and open it with SSMS.
2. Modify the following parameters as needed:

Configuration item	Description
@backup_dat abases_list	Databases to be backed up. Separate multiple databases with semicolons (;) or commas (,).
@backup_type	Backup type. Values are as follows: · FULL: full backup · DIFF: differential backup · LOG: log backup
@backup_folder:	Local folder that stores the backup file. It will be automatica lly created if it does not exist.
@is_run	Whether to perform a backup. Values are as follows: · 1: Perform a backup. · 0: Only perform a check.

3. Run the backup script.

Upload the backup file to OSS

Use any of the following methods to upload the backup file to your OSS bucket:

Method 1: Use ossbrowser

It is recommended that you use the ossbrowser tool to upload the backup file to OSS. For more information, see *ossbrowser*.

Method 2: Use the OSS console

If the backup file is smaller than 5 GB, you can use the OSS console to upload it. For more information, see *Upload an object*.

Method 3: Use an OSS API

If you require automatic migration, use an OSS API to perform an upload that can be paused and resumed. For more information, see *Multipart upload*.

Create a migration task

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the target instance ID to enter the Basic Information page.
4. On the left-side navigation pane, click Backup and Recovery.
5. Click OSS Backup Data Upload at the upper right corner.

6. If you are using the function for the first time, authorize the RDS official service account to access OSS:

   a. In the import data step of the Import Guide, click Authorize.

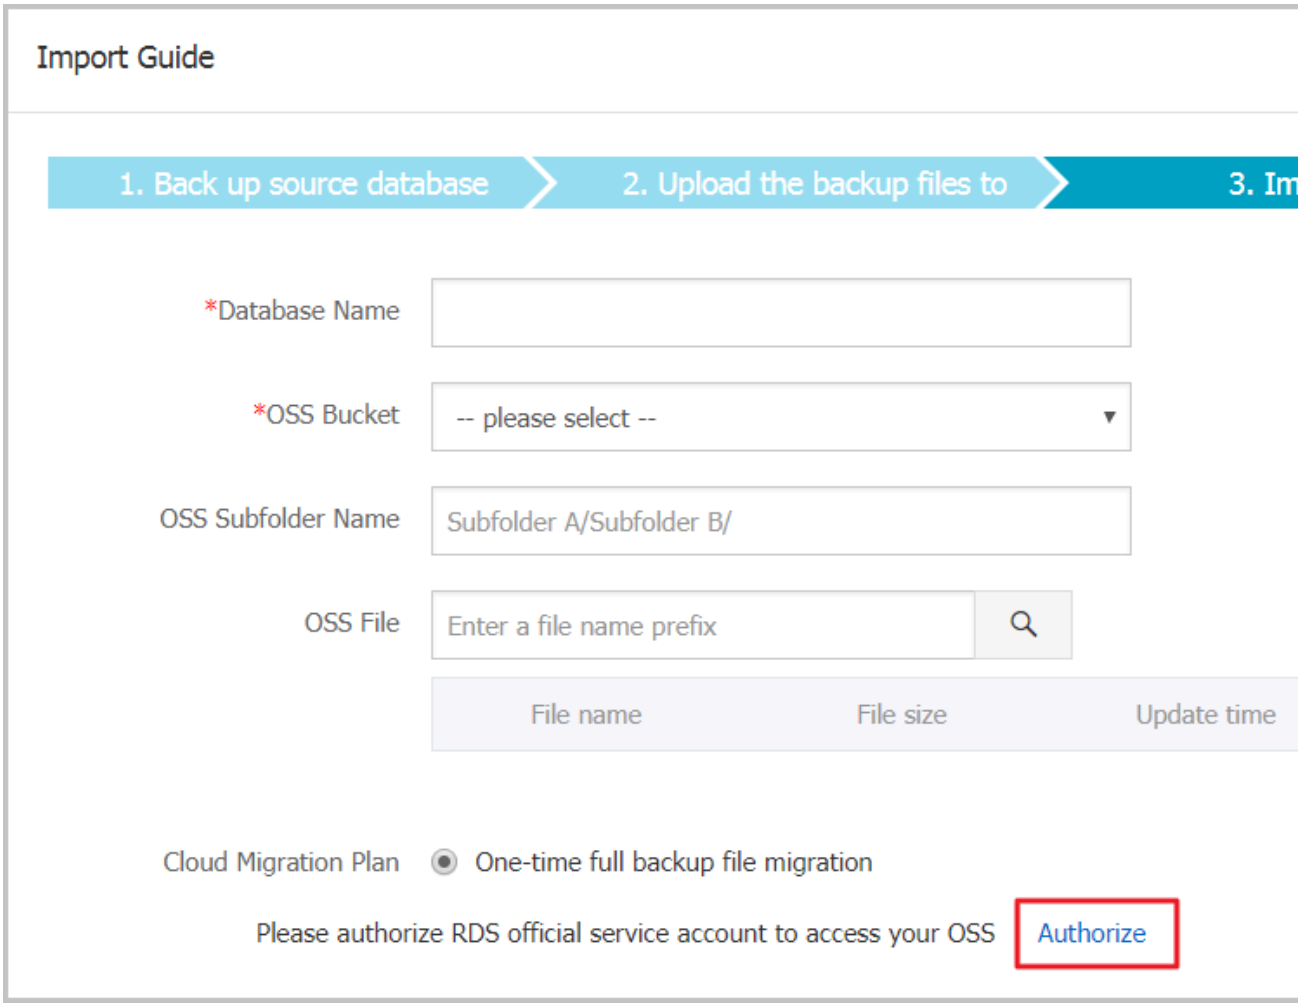

   b. Click Confirm Authorization Policy.

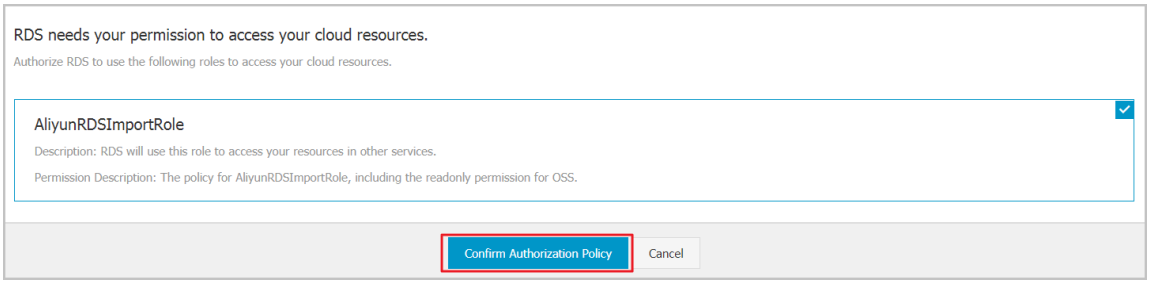

7. Set the following parameters and click OK to generate an OSS backup file upload task.

Configuration item	Description
Database Name	Target database name in the target instance
OSS Bucket	OSS bucket that stores the backup file
OSS Subfolder Name	Name of the subfolder where the backup is located.
OSS File	Click the magnifier icon on the right. You can perform a fuzzy search with the backup file prefix. The file names, sizes, and update time are displayed. Select the backup file you need.
Cloud Migration Plan	· Immediate Access (Full Backup): If you have only the full backup file, select Immediate Access. · Access Pending (Incremental Backup): If you have a full backup file and a differential or log backup file, select this option.
Consistency Check Mode	· Synchronous DBCC: Perform DBCC check only after the database is opened. This reduces service downtime because DBCC check takes a long time if the database is large. If you are sensitive to service downtime and do not care about the DBCC check result, select this option. · Asynchronous DBCC: If you want to use DBCC check to find out consistency errors of your source database, select this option. Note that this option lengthens the time it takes to open the database.

You can click Refresh to view the latest status of the migration task. If the migration fails, view the task description and rectify faults by referring to Common errors at the end of this document.

View migration records

View migration records as follows:

On the Backup and Recovery page, click Backup Data Upload History. Migration records of the past week are displayed by default. You can change the query time range as needed.

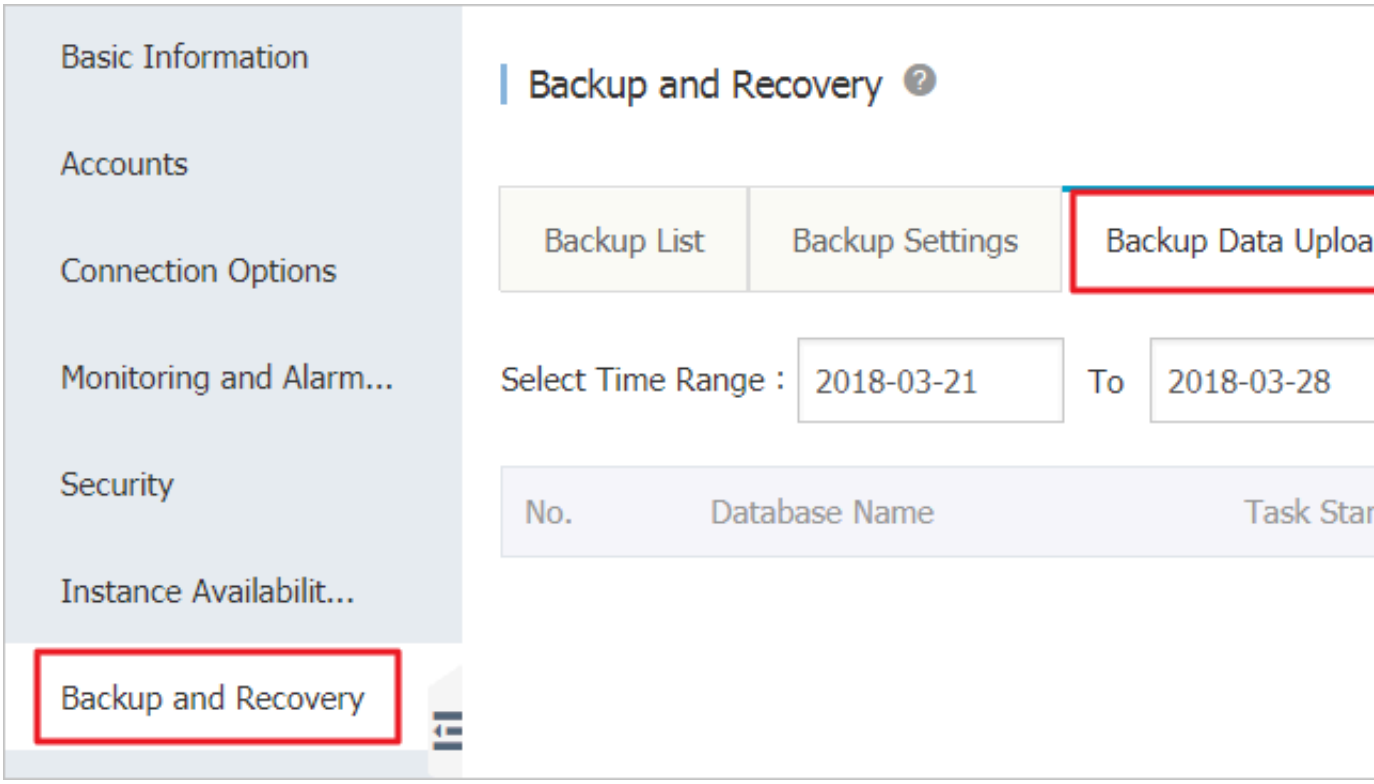

Common errors

Each migration record has a task description, which helps you identify the failure cause. Common errors are as follows:

Database with the same name already exists

- Error message: The database (xxx) is already exist on RDS, please backup and drop it, then try again.
- Error cause: An existing database with the same name is not allowed in the target instance. This prevents you from mistakenly overwriting a database.
- Solution: If a database with the same name already exists in the target instance , perform a full backup of the database on the console and delete the database before the migration.

Differential backup files

- Error message: Backup set (xxx.bak) is a Database Differential backup, we only accept a FULL Backup.
- Error cause: The migration supports only full backup files rather than differential backup files.

Transaction log backup files

- Error message: Backup set (xxx.trn) is a Transaction Log backup, we only accept a FULL Backup.
- Error cause: Full migration supports only full backup files rather than log backup files.

Backup file verification fails

- Error message: Failed to verify xxx.bak, backup file was corrupted or newer edition than RDS.
- Error cause: The verification fails because the backup file is damaged or the local SQL Server version is later than the target RDS SQL Server version. For example, the verification fails if the migration is from SQL Server 2016 to SQL Server 2012.
- Solution: If the backup file is damaged, perform a full backup again to generate a new backup file. If the local SQL Server version is later than the target RDS SQL Server version, change the target RDS SQL Server version.

DBCC CHECKDB errors

- Error message: DBCC checkdb failed
- Error cause: DBCC CheckDB failure indicates that the local database has errors.
- Solution:

1. Run the following command to fix the local database (this may cause data loss):

```
DBCC   CHECKDB  ( DBName ,  REPAIR_ALL  OW_DATA_LO  SS )  WITH
NO_INFOMSG  S ,  ALL_ERRORM  SGS
```

2. Perform a full backup for the database again.
3. Upload the new database file to OSS.
4. Perform the migration again on the RDS console.

OSS download link expires

This error only happens to the RDS for SQL 2008 R2 High-availability Edition instances.

- Error message: Failed to download backup file since OSS URL was expired.
- Error cause: The OSS download link has expired, so the backup file download fails.

- Solutions:
  - Solution 1: Set the download link validity period to a larger value (at most 18 hours).

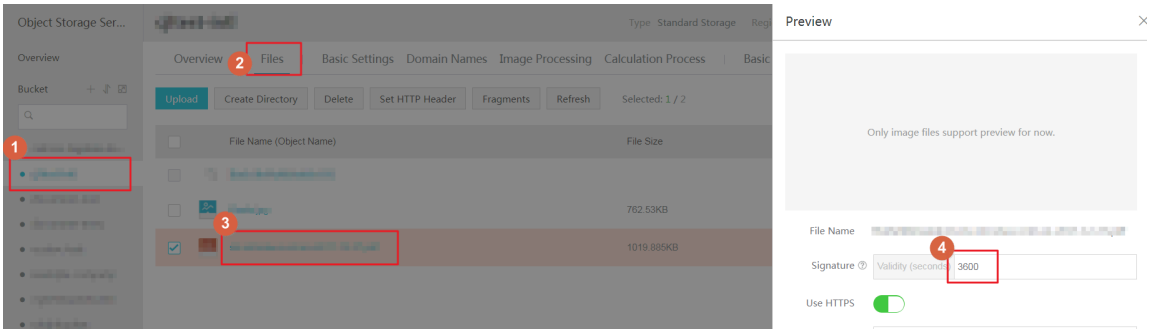

  - Solution 2: Set the ACL permission of the OSS database backup file to Public Read.

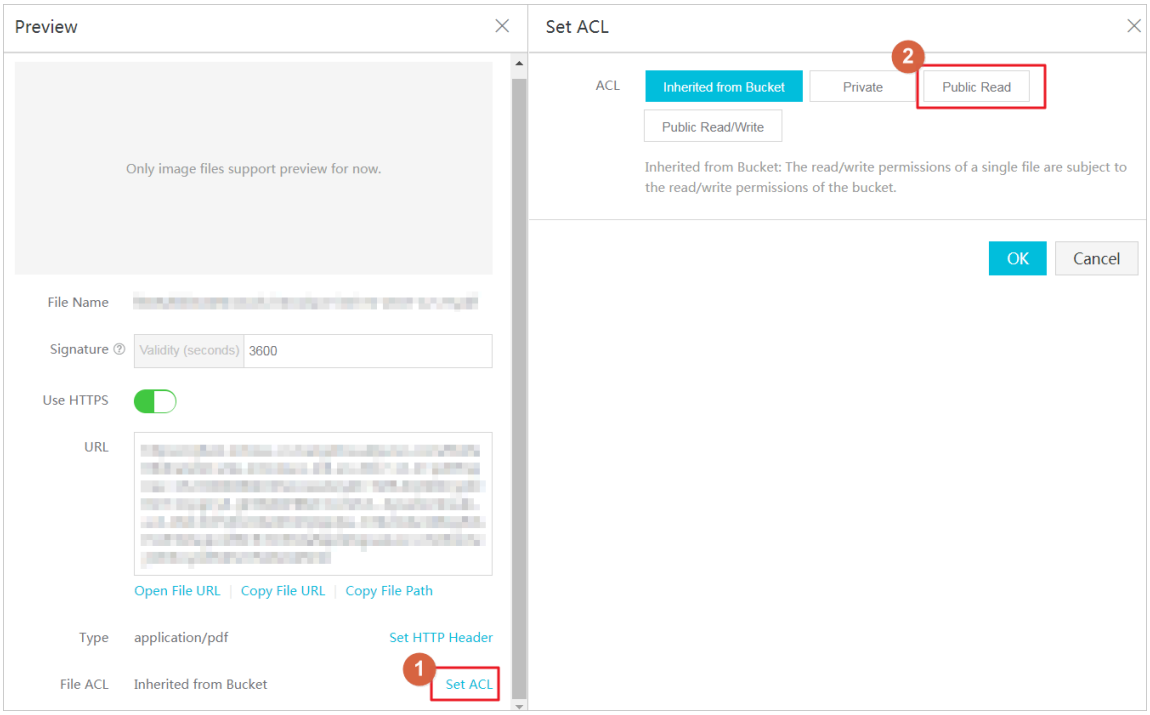

> **Note:**
> The backup file with the Public Read ACL permission can always be downloaded without an expiration date. To prevent security risks, set the ACL permission to Private after migrating the file.

Insufficient space 1

- Error message: Not Enough Disk Space for restoring, space left (xxx MB) < needed (xxx MB)
- Error cause: The remaining space on the instance is insufficient for migration.

- Solution: Expand the storage space of the instance.

Insufficient space 2

- Error message: Not Enough Disk Space, space left xxx MB < bak file xxx MB
- Error cause: The remaining space on the instance is smaller than the backup file size.
- Solution: Expand the storage space of the instance.

No superuser account

- Error message: Your RDS doesn’t have any init account yet, please create one and grant permissions on RDS console to this migrated database (XXX).
- Error cause: If the RDS instance has no superuser account, the migration still succeeds, but the migration task does not know which user to authorize.
- Solution:
  1. Create a superuser account. For details, see *Create accounts and databases (SQL Server 2012 or 2016)*.
  2. Reset the password of the superuser account. For more information, see *Reset the instance password*.
  3. Use the superuser account to access the database on the cloud.

3.7 Migrate a MySQL database from Tencent Cloud to Alibaba Cloud

This topic describes how to migrate a MySQL database from Tencent Cloud to Alibaba Cloud and the corresponding precautions.

Prerequisites

- You have *created an RDS instance*.
- You have *created an account with read and write permissions*.

Limits

- Structure migration does not support migration of events.
- For MySQL databases, DTS reads floating-point values (FLOAT and DOUBLE data types) with `round ( column , precision )`. If the column definition does not specify the precision, the precision is 38 for FLOAT values and 308 for DOUBLE values.

- If the object name mapping function is used for an object, migration of objects relying on the object may fail.
- For incremental migration, you need to enable binary logging for the MySQL instance in the source database.
- For incremental migration, binlog_format of the source database must be set to ROW.

> **Note:**
> You can modify parameters of Tencent Cloud databases by choosing Manage Database > Parameter Settings.

-
- For incremental migration, if the source instance has binlog file ID disorder caused by cross-host migration, the incremental migration may have data loss.

Precaution

DTS automatically attempts to recover abnormal tasks of the past seven days. This may cause the new data in the target instance to be overwritten by the source database data. Therefore, you must revoke the write permission of the DTS account that is used to access the target instance by running the `revoke` command.

Procedure

1. Log on to your MySQL database instance on Tencent Cloud. On the Instance Details page, view the details of `Public IP` , including the domain name and port.

> **Note:**
> If an Internal IP address is not enabled, you need to click Enable, and then click OK in the displayed dialog box.

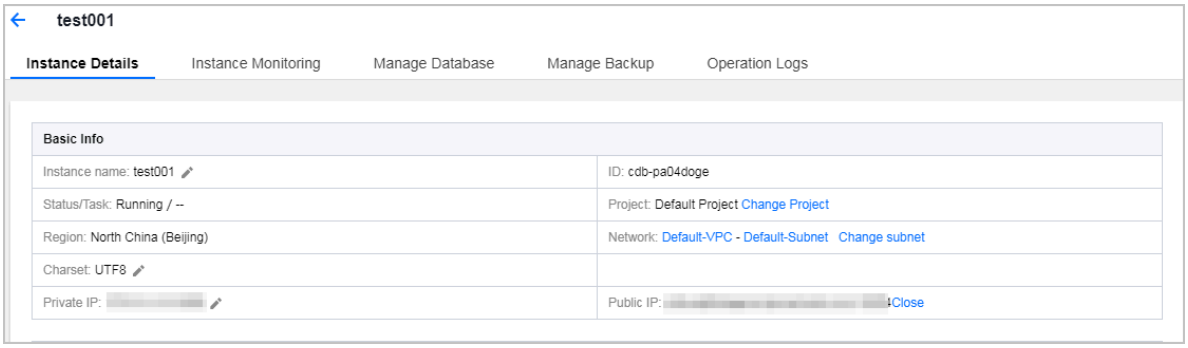

2. Log on to the *DTS Console*.

3. In the left-side navigation pane, click Data Migration. In the right pane, click Create Migration Task in the upper-right corner.
4. Enter information about the source and target databases. The following table describes the parameters.

Database type	Parameter	Description
Source database (on Tencent Cloud)	Instance Type	Type of the instance in the source database. Select `On-premises Databases`.
	Instance Region	If you have configured access control for your instance, you need to allow the specified Internet IP segment of the region to access the instance before configuring a migration task. Note: You can click Get DTS IP to view and copy the IP segment of the region.
	Database Engine	Source database type. Select `MySQL`.
	Host Name or IP Address	Domain name in `Public IP`
	Port	Port in `Public IP`
	Database account	Default superuser account `root`
	Database Password	Password of the root account
Target database (on Alibaba Cloud)	Instance Type	Type of the instance in the target database. Select `RDS Instance`.
	Instance Region	Region of the target instance
	RDS Instance ID	ID of the instance in the selected region. Select the ID of the target instance.
	Database account	An account with read and write permissions under the target instance
	Database Password	Account password

Database type	Parameter	Description
	Connection method	Select Non-encrypted connection or SSL secure connection. The latter greatly increases CPU consumption.

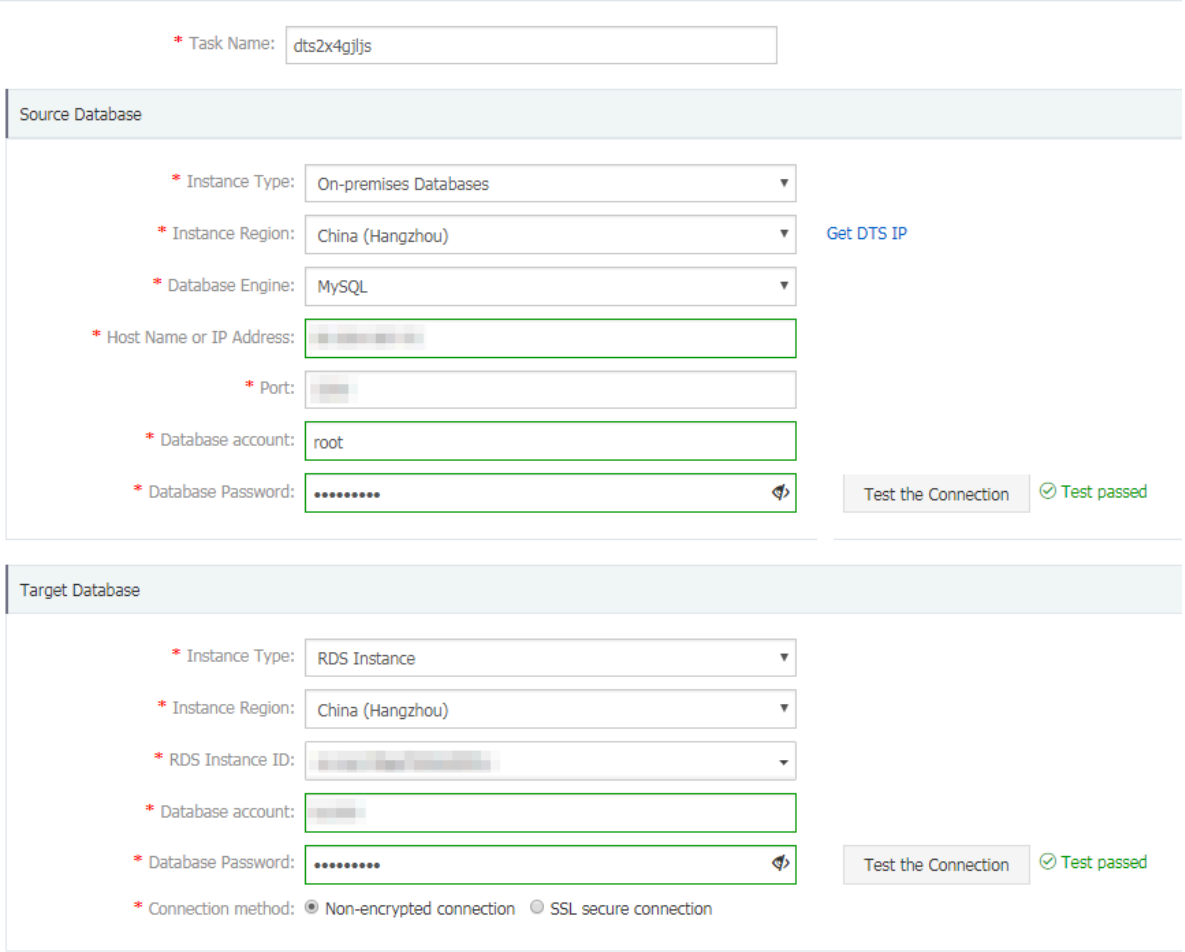

5. Click Test the Connection and confirm that the test results for both the source and target databases are `Test passed`.
6. Click Authorize Whitelist and Enter into Next Step.
7. Select the migration type. In the Migration objects area, select the target database and click to add the database to the Selected objects area.

Note:

To maintain data consistency before and after migration, we recommend that you migrate the structure, full data, and incremental data.

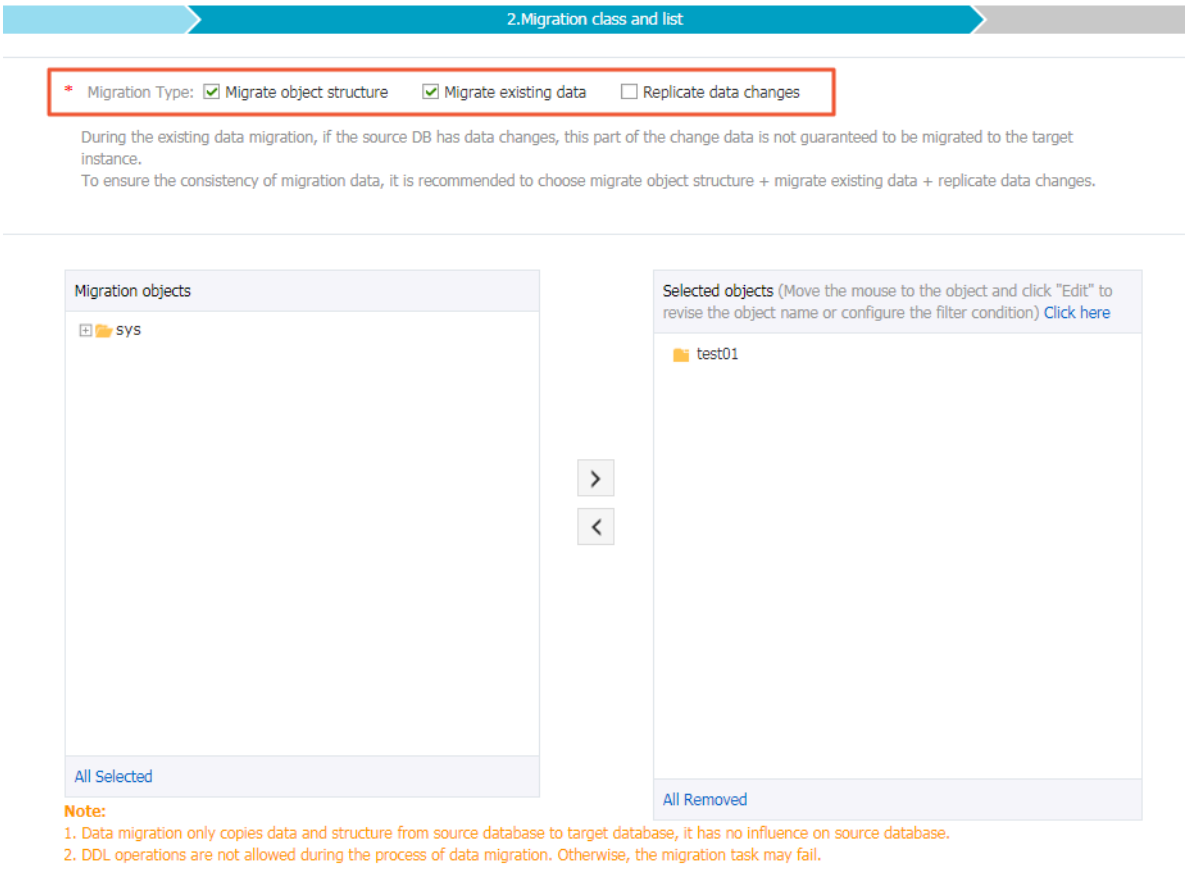

8. Click Pre-check and Start and wait until the pre-check ends.

Note:

If the check fails, you can rectify faults according to error items and restart the task.

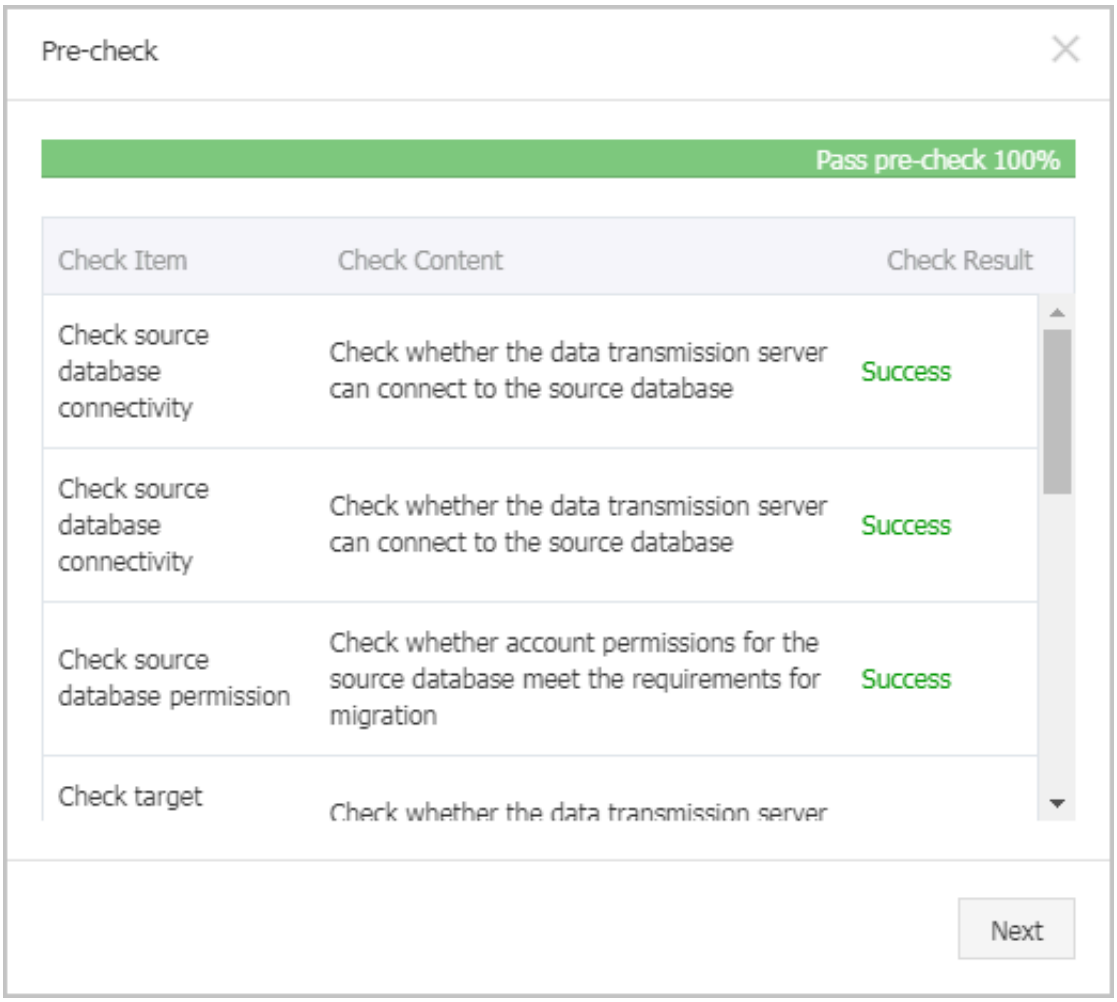

9. Click Next. In the Confirm Purchase Configuration dialog box, read and confirm you agree to the Service Terms of Data Transmission (Pay-As-You-Go) and click Buy and Start Now.

Note:
Currently, structure migration and full migration are free of charge, while incremental migration is charged by the hour according to link specifications.

10.Wait until the migration task is completed.

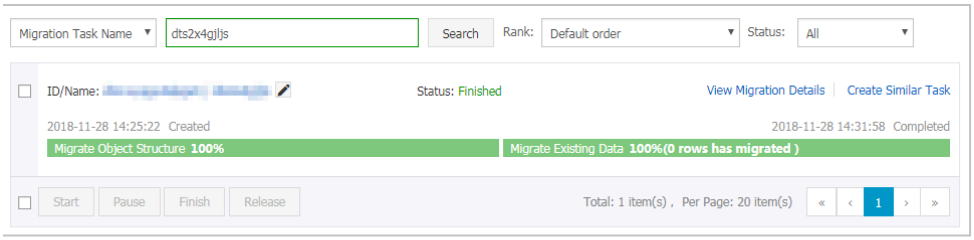

3.8 Migrate a MySQL database from Google Cloud to Alibaba Cloud

This topic describes how to migrate a MySQL database from Google Cloud to Alibaba Cloud and the corresponding precautions.

Prerequisites

- You have *created an Alibaba Cloud RDS MySQL instance*.
- You have *created an account with read/write privileges*.

Limits

- Structure migration does not support migration of events.
- For MySQL databases, DTS reads floating-point values (FLOAT and DOUBLE data types) with `round ( column , precision )`. If the column definition does not specify the precision, the precision is 38 for FLOAT values and 308 for DOUBLE values.
- If the object name mapping function is used for an object, migration of objects relying on the object may fail.
- For incremental migration, you need to enable binlog for the source MySQL instance.
- For incremental migration, binlog_format of the source database must be set to ROW.

  Note:
  You can modify parameters of Google Cloud databases by choosingInstance details > Configuration > Edit configuration > Add database flags.
- For incremental migration, if the source database version is MySQL 5.6 or later, binlog_row_image must be set to FULL.

- For incremental migration, if the source instance has binlog file ID disorder caused by cross-host migration, the incremental migration may have data loss.

Precautions

DTS automatically attempts to recover abnormal tasks of the past seven days. This may cause the new data in the target instance to be overwritten by the source database data. Therefore, you must revoke the write permission of the DTS account that is used to access the target instance by running the `revoke` command.

Procedure

1. Log on to your database instance on Google Cloud. On the Instance details page, view `Public IP address`.

> **Note:**
> If an Internal IP address is not enabled, perform related settings by going to Configuration > Edit configuration > Set connectivity.

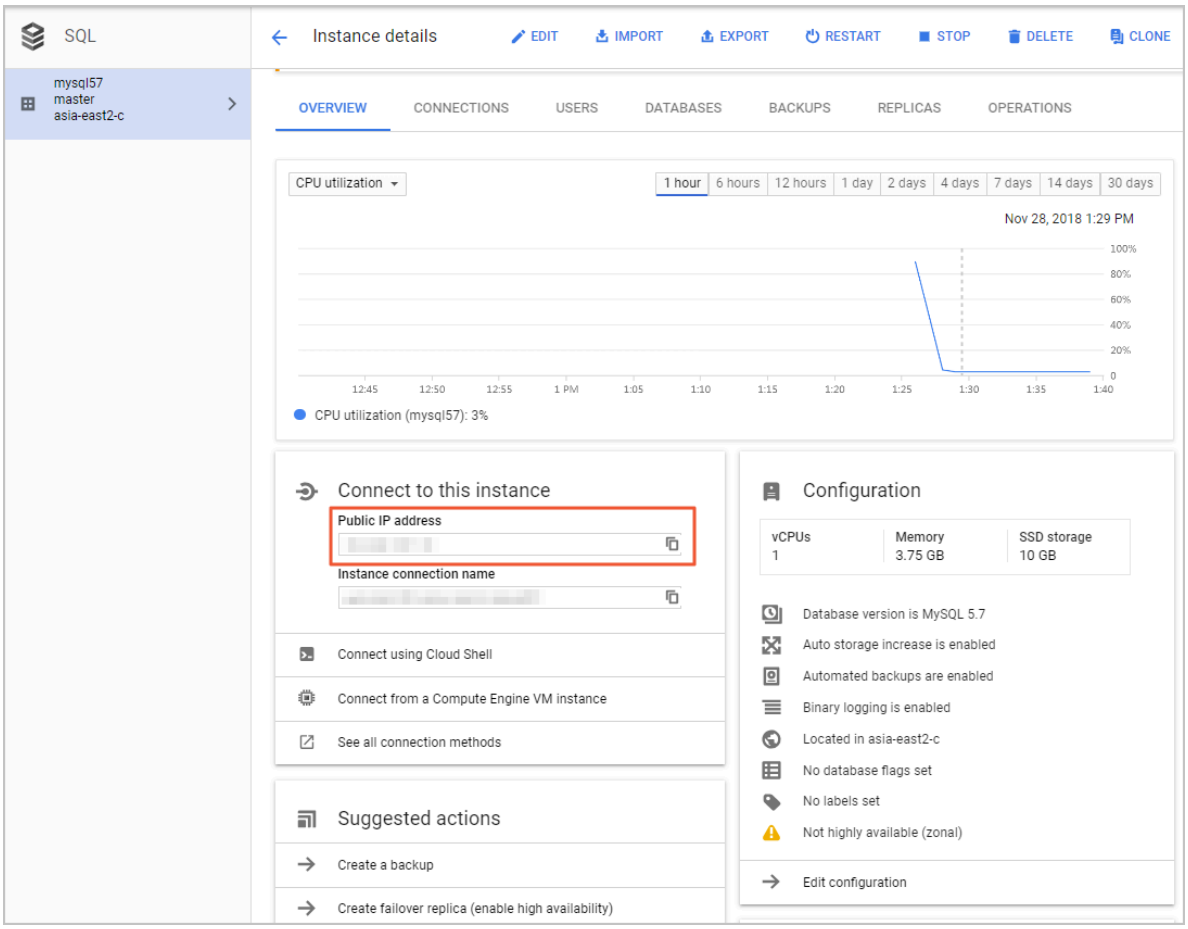

2. Choose Configuration > Edit configuration > Set connectivity > Add network, and then add the IP address of the *region of the source database instance* obtained from DTS.

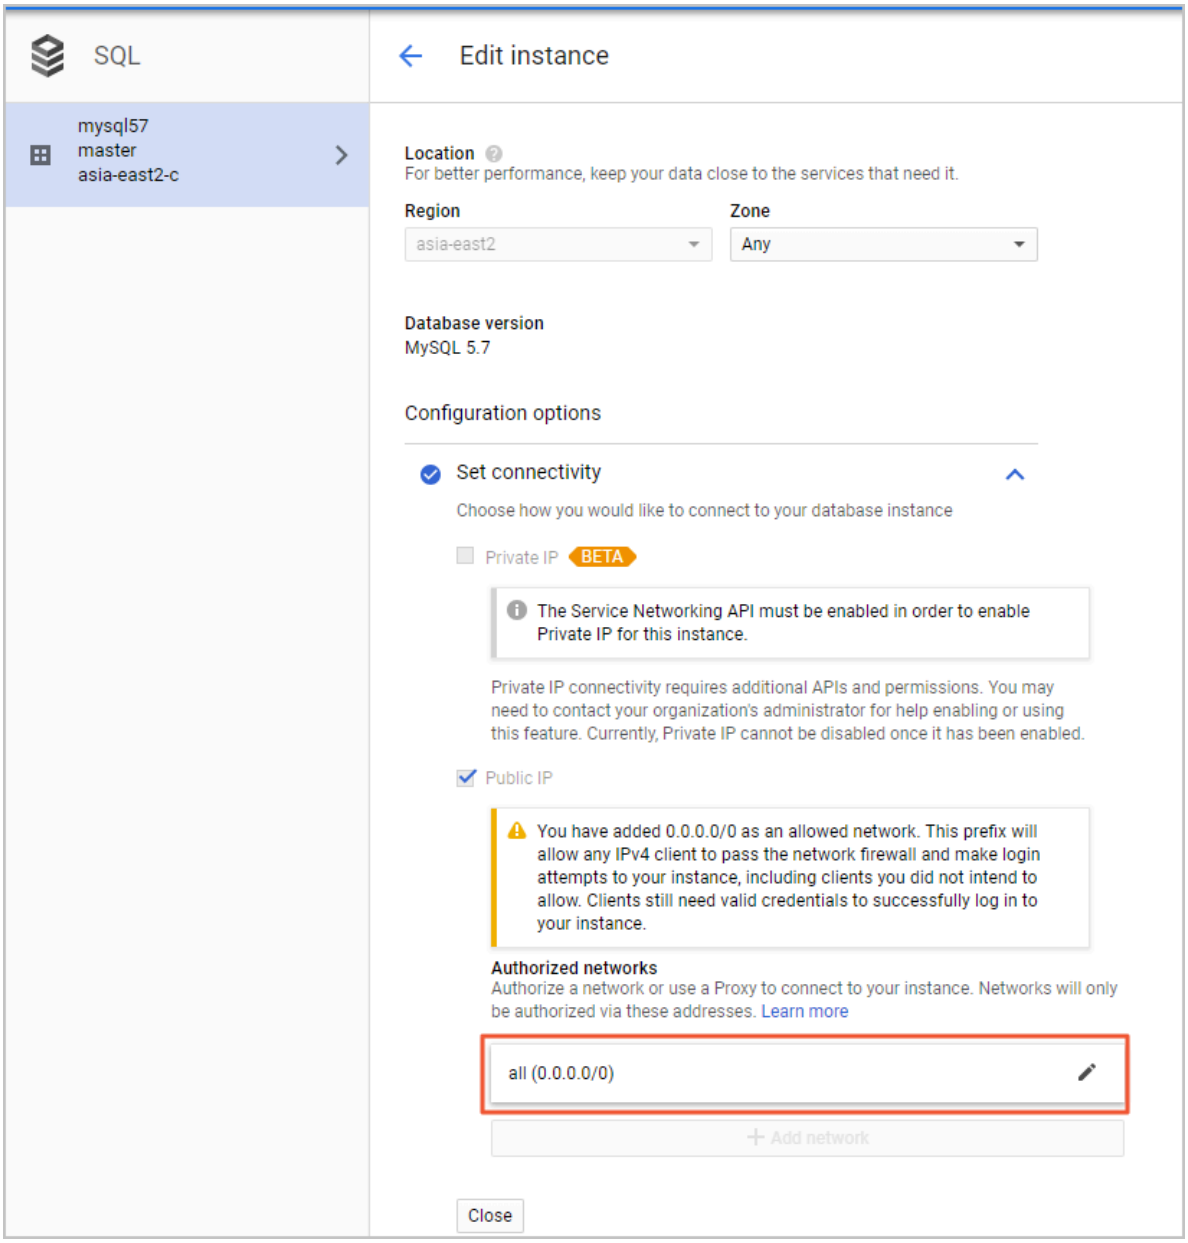

3. Log on to the *DTS Console*.
4. In the left-side navigation pane, click Data Migration. In the right pane, click Create Migration Task in the upper-right corner.

5. Enter information about the source and target databases. The following table describes the parameters.

Database type	Parameter	Description
Source database (on Google Cloud)	Instance Type	Type of the instance in the source database. Select `On-premises Databases`.
	Instance Region	If you have configured access control for your instance, you need to allow the specified Internet IP segment of the region to access the instance before configuring a migration task. Note: You can click Get DTS IP to view and copy the IP segment of the region.
	Database Engine	Source database type. Select `MySQL`.
	Host Name or IP Address	`Public IP address` of the database
	Port	Default port 3306
	Database account	Default superuser account `root`
	Database Password	Password of the root account
Target database (on Alibaba Cloud)	Instance Type	Type of the instance in the target database. Select `RDS Instance`.
	Instance Region	Region of the target instance
	RDS Instance ID	ID of the instance in the selected region. Select the ID of the target instance.
	Database account	An account with read and write permissions under the target instance
	Database Password	Account password

Database type	Parameter	Description
	Connection method	Select Non-encrypted connection or SSL secure connection. The latter greatly increases CPU consumption.

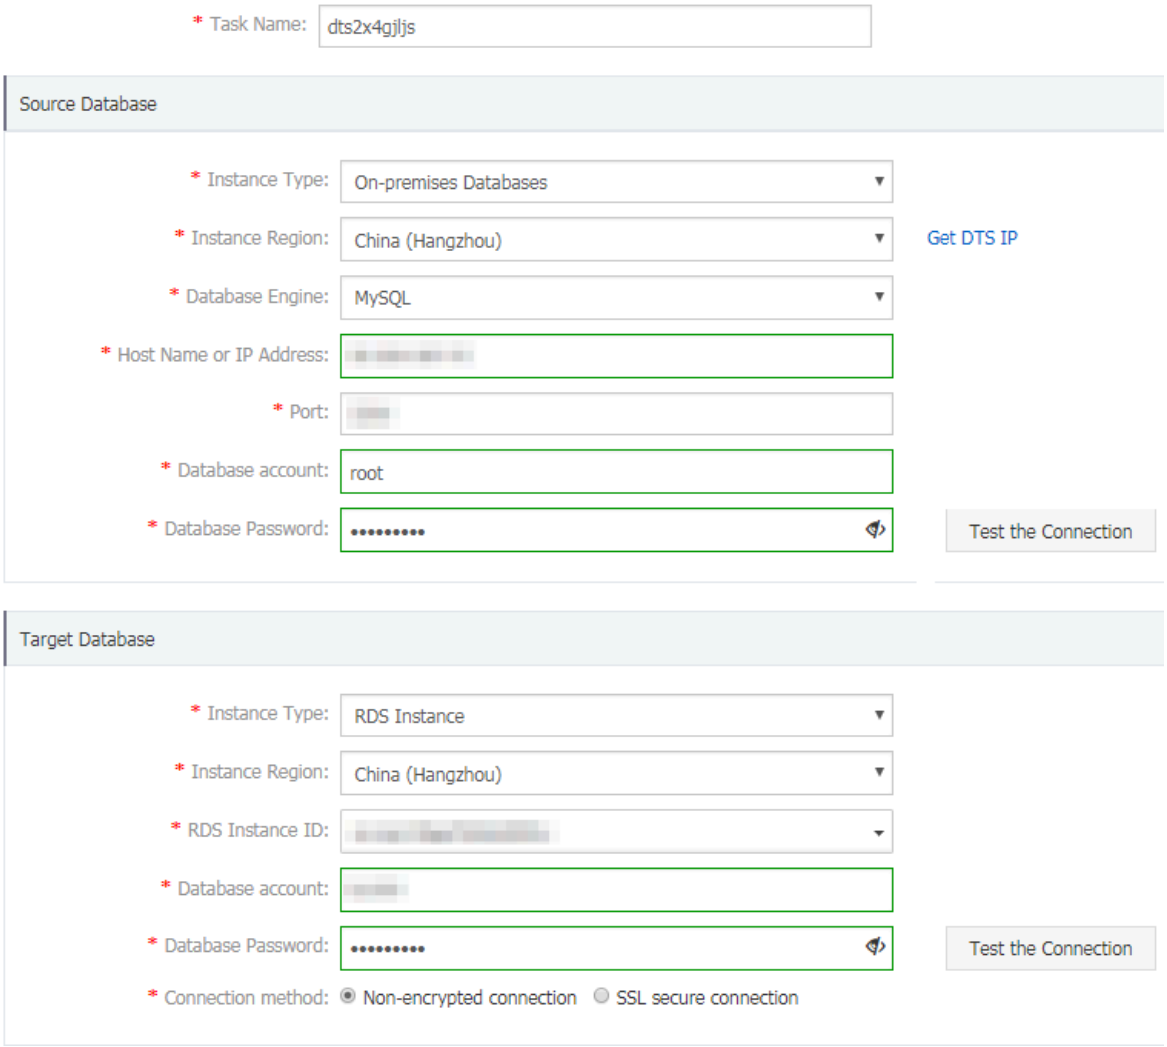

6. Click Test the Connection and confirm that the test results for both the source and target databases are `Test passed` .
7. Click Authorize Whitelist and Enter into Next Step .
8. Select the migration type. In the Migration objects area, select the target database and click to add the database to the Selected objects area.

Note:

To maintain data consistency before and after migration, we recommend that you migrate the structure, full data, and incremental data.

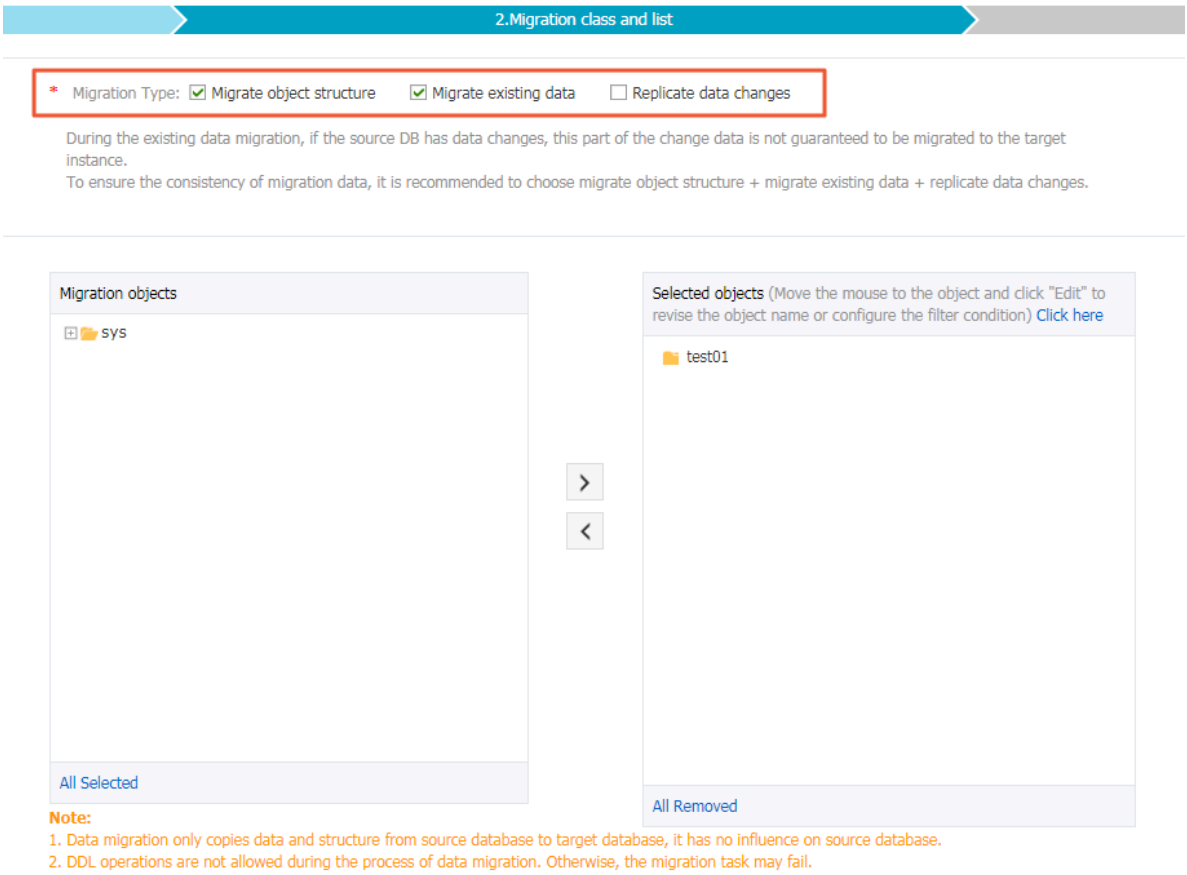

9. Click Pre-check and Start and wait until the pre-check ends.

Note:

If the check fails, you can rectify faults according to error items and restart the task.

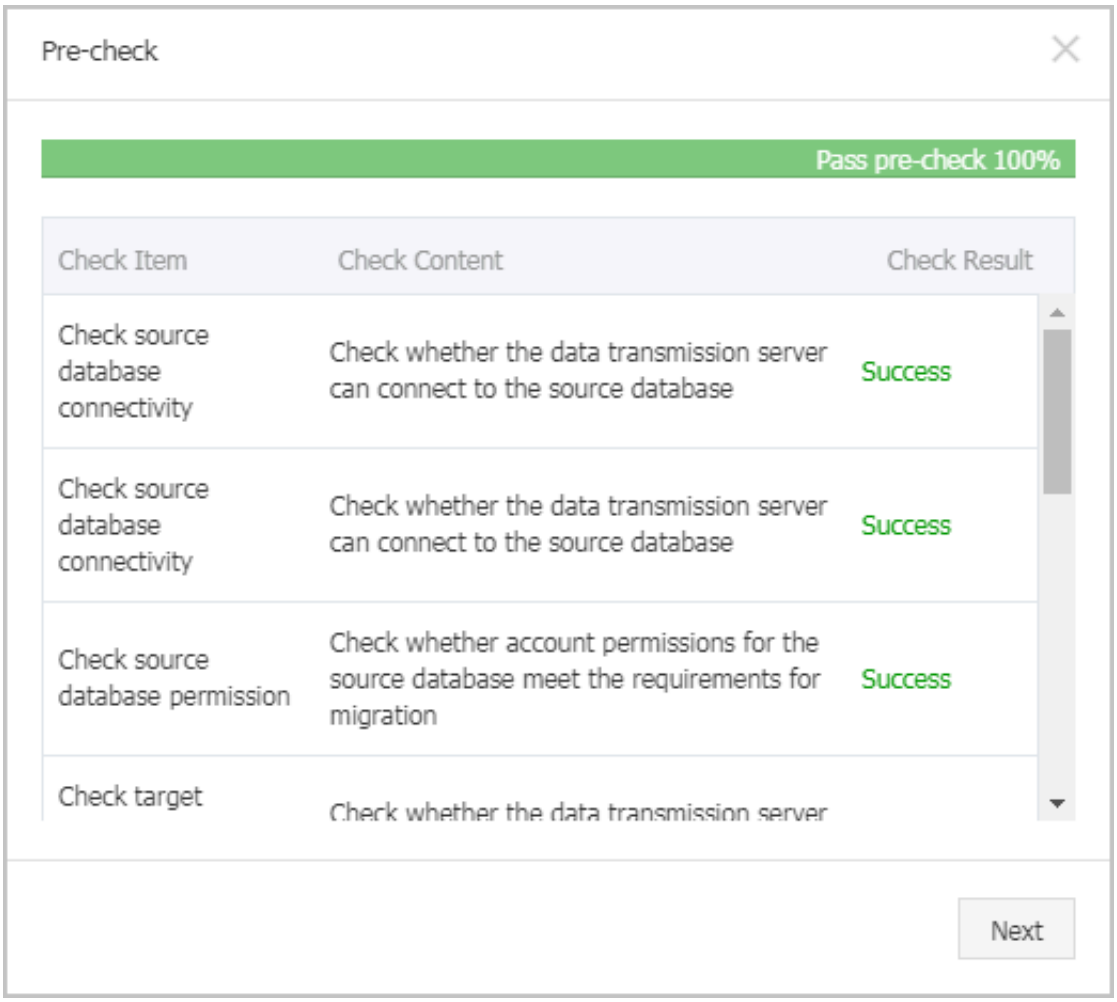

10. Click Next. In the Confirm Purchase Configuration dialog box, read and confirm you agree to the Service Terms of Data Transmission (Pay-As-You-Go) and click Buy and Start Now.

Note:
Currently, structure migration and full migration are free of charge, while incremental migration is charged by the hour according to link specifications.

11.Wait until the migration task is completed.

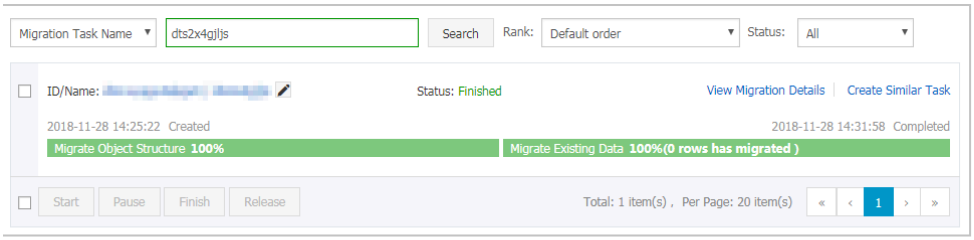

4 Billing management

4.1 Change the billing method

You can change a Pay-As-You-Go instance to a Subscription instance.

Attention

- Think twice before such a conversion, because a Subscription instance cannot be converted back to a Pay-As-You-Go instance.
- Within the contract period of a Subscription instance, you can only upgrade it but cannot downgrade or release it.
- After the conversion is successful, the Subscription billing method is immediately applied. For more information, see *Pricing*.
- An order is generated when you change a Pay-As-You-Go instance to a Subscription instance. The conversion takes effect only after you pay for the order. If you leave the order unpaid, the order is displayed on the *Orders* page and you cannot purchase new instances or change billing methods of instances.

  Note:

  - If you upgrade an instance when its billing method change order is unpaid, you cannot pay for the order any more because the order amount is insufficient. Invalidate the order and change the billing method again.
  - If you do not want to pay for an order, invalidate it on the *Orders* page.

Prerequisites

- You are the owner of the instance.
- The instance type is not a history instance type. For more information, see *Instance type overview*.

  Note:

  A Pay-As-You-Go instance of a history type cannot be converted to a Subscription instance. To change the billing method for a Pay-As-You-Go instance of a history type, change the instance type to a new type first. For operation details, see *Change configurations*.

- The billing method of the instance is Pay-As-You-Go, and the instance status is Running.

  > **Note:**
  > After you submit the order, if the instance status changes (for example, to the Locked state), payment will fail. You can pay for the order only when the instance status restores to Running.

- There is no unfilled billing method change order (namely, new Subscription instance order) of an instance .

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the instance ID to enter the Basic Information page.
4. In the Status area, click Subscription Billing.

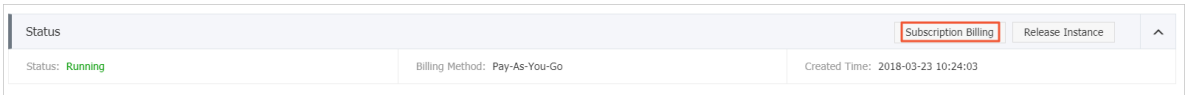

5. Select the subscription period.
6. Click Pay Now and pay for the order.

4.2 Enable auto-renewal for a Subscription instance

Auto-renewal for a Subscription instance frees you from regular manual renewals. It also avoids service interruptions caused if the instance expires and is not renewed in time.

If you did not select auto-renewal when you purchased the Subscription instance, you can set it up on the Alibaba Cloud Billing Management console. When the setup is done, the subscription is automatically renewed based on the selected renewal cycle. For example, if you select a three-month renewal cycle, three months of subscription is automatically paid for each renewal. This document explains how to enable auto-renewal for your Subscription instance.

Prerequisite

You have logged on to Alibaba Cloud console with your master account.

Attentions

- The renewal cycle cannot be changed while enabling the auto-renewal function. For variable renewal cycles, renew the instance manually. For more information about how to handle manual renewal, see *Manually renew a Subscription instance*.
- If you select auto-renewal, you are charged three days before the instance expires. Credit cards and coupons are supported for each renewal payment.
- If you manually renew your instance before the charging date, auto-renewal takes place based on the new expiration date.
- The auto-renewal function takes effect the next day after it is enabled. If your instance expires on the next day, manually renew it to prevent service interrupti ons.

Procedure

1. Log on to the *Billing Management* console of Alibaba Cloud.
2. In the left-side navigation pane, select Renewal.
3. Select ApsaraDB for RDS in the Product drop-down list, and select the region where the target instance is located and its creation date. Alternatively, select the default search range.
4. Click Search.

5. In the Auto-renewal column for the target instance, move the slider to the right.
6. On the open automatic page, set automatic renewal hours.
7. Click Open automatic.

4.3 Manually renew a Subscription instance

A Subscription instance must be renewed within 15 days after expiration. Subscription instances are automatically released when the payment is overdue for 15 days. As a result, all data for the instance is deleted and cannot be recovered. For more information about renewal, see *Renewal*.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.

3. Click the ID of the target instance to go to the Basic Information page.
4. Click Renew in Status area, as shown in the following figure.

5. Select the renewal period on the Renew page.

> Note:
> You can change the configuration if needed.

6. Read and confirm the terms of service, then select I agree to Product Terms of Service and Service Level Notice.
7. Click Pay to complete the payment process.

Related topic

Enable auto-renewal of the subscription instance

5 Instance management

5.1 Restart an instance

Context

You can manually restart an instance when the number of connections exceeds the threshold or any performance issue occurs for the instance. Restarting an instance may interrupt connections. Proceed with caution and make appropriate service arrangements before restarting an instance.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance or click Manage to enter the Basic Information page.
4. Click Restart Instance in the upper right corner on the instance management page. In the displayed dialog box, click OK.

5.2 Configure the maintenance period

RDS needs to be regularly maintained to guarantee overall instance health in production environment. You can set the maintenance period in the idle service hours based on service regularities to prevent potential interruptions for production during maintenance. RDS performs regular maintenance operations during the maintenance period you have configured.

Background information

To guarantee stability and efficiency of ApsaraDB RDS instances on the Alibaba Cloud platform, the backend system performs a serial of maintenance tasks at an irregular basis as needed.

Before official maintenance, RDS sends text messages and emails to contacts configured by your Alibaba Cloud account.

To guarantee stability during the maintenance process, instances enter the Instance being maintained state before the preset maintenance period on the day of

maintenance. When an instance is in this state, normal data access to databases is not affected. However, apart from account management, database management, and IP address addition to the whitelist, other services involving changes (such as common operations including upgrade, degrade, and restart of the instance) are unavailable on the console. Query services such as performance monitoring are available.

When the maintenance period begins, transient disconnection occurs once or twice to the instance during this period. Make sure that applications support the reconnecti on policy so that the instance can be restored to the normal state after transient disconnection.

Procedure

1. Log on to the *RDS console* and select the target instance.
2. Select Basic information in the menu.
3. In the Configuration information area, click Settings following Time segment. The default maintenance period of RDS is from 02:00 to 06:00.

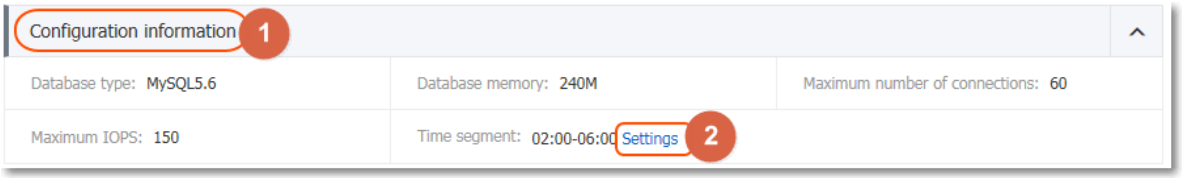

4. Select the maintenance period and click Save, as shown in the following figure.

Note:
Note: The ime segment is tn Beijing Time.

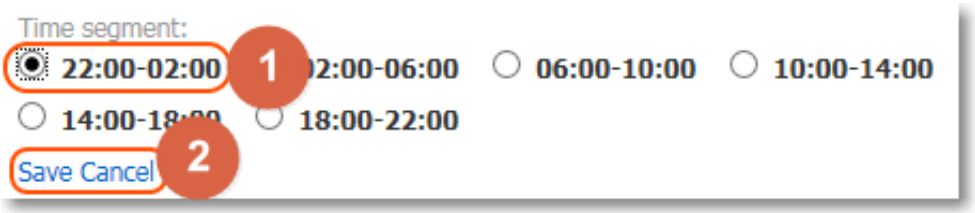

5.3 Migrate instance across zones

If the zone in which the instance is located is in full load or the instance performance is affected for other reasons, you can migrate the instance to other zones in the same region. During the migration, the RDS service is interrupted and certain operations cannot be performed. Therefore, we recommend that you set the migration time to off -peak hours. This document describes migration details.

Note:
Currently, only MySQL 5.5/5.6, SQL Server 2008 R2, PostgreSQL 9.4, PPAS 9.3 instances support instance migration across zones.

Background information

You can select between single-zone and multi-zone instances. A multi-zone is a physical area created through combination of multiple single zones in the same region. For example, you can create multi-zone 1 by combining zone B and zone C. Compared to single-zone instances, multi-zone instances can withstand high-level disasters. For example, single-zone instances can withstand faults at the server and rack level, while multi-zone instances can withstand faults at the data center level.

Currently, multi-zones are supported in China East 1 (Hangzhou), China East 2 (Shanghai), China North 2 (Beijing), China South 1 (Shenzhen), Hong Kong, and Singapore (the regions supporting multi-zones may be updated. Select one of the available options on the RDS console). No extra fee is charged for the use of a multi-zone.

If the zone in which the instance is located is in full load or the instance performance is affected for other reasons, you can migrate the instance to other zones in the same region. Instance migration across zones involves copying the instance data to the new zone, and the migration is performed at the instance level. After the instance is migrated to a new zone, all its attributes and configurations remain the same. It often takes several hours to migrate an instance to a new zone, and the time is subject to the instance size. After all the instance data is copied to the new zone, the instance is deleted from the original zone.

You can choose one of the following methods to migrate an instance across zones:

- Migrate the instance from a single zone to another single zone.
- Migrate the instance from a single zone to a multi-zone. In this case, if the instance has a master database and a slave database, the two databases are randomly allocated in the multi-zone. For example, when an instance with a master database and a slave database is migrated from zone A to multi-zone 1 (zone B + zone C), if the master database is allocated to zone B, the slave database is allocated to zone C.
- Migrate the instance from a multi-zone to a single zone. In this case, the master and slave databases of the instance are migrated to the same zone, and the instance can withstand lower-level disasters.

Note:

Because certain network delay exists between multi-zones, the response time of a multi-zone instance to a single update may be longer than that of a single-zone instance when a multi-zone instance adopts the semi-synchronous data replication mode. In this case, increase the overall throughput by enhancing the concurrency.

Attentions

- **Migration across zones is possible only when the region of an instance has multiple zones.**
- **During the migration across zones, most management operations cannot be performed. Therefore, choose an appropriate time for the migration. The following lists the operations that can or cannot be performed:**

Operation	Whether the operation can be performed
Modify the whitelist	Yes
Enable SQL audit	Yes
Set the maintenance period	Yes
Add read-only instances	No
Add disaster recovery instances	No
Release an instance	No
Change the billing method to the Subscription mode	No
Change configurations	No
Create a common or master account	No
Reset the account password	No
Modify account permissions	No
Create and delete databases	No
Change the network type	No
Change the access mode	No
Modify the connection address	No
Apply for an Internet IP address	No

Operation	Whether the operation can be performed
Switch between master and slave databases	No
Change the data backup mode	No
Restore instance data	No
Modify parameters	No

- There is a 30 seconds of transient disconnection during migration across zones. Make sure that your application has a reconnection policy.

Procedure

1. Log on to the *RDS console*.
2. Select the region of the target instance.
3. Click the target instance ID to go to the Basic Information page.
4. Click Migration Across Zones in the Basic Information area, as shown in the following figure.

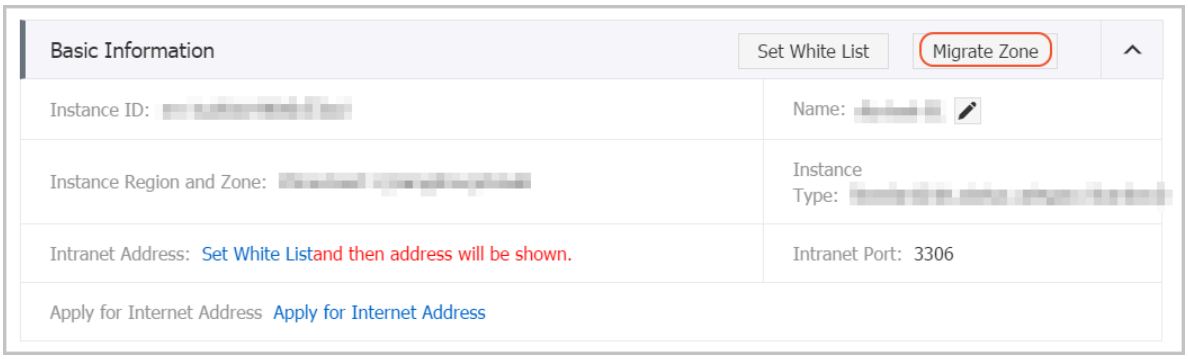

5. Select a target zone in the Migrate Instance to Other Availability Zones dialog box, as shown in the following figure.

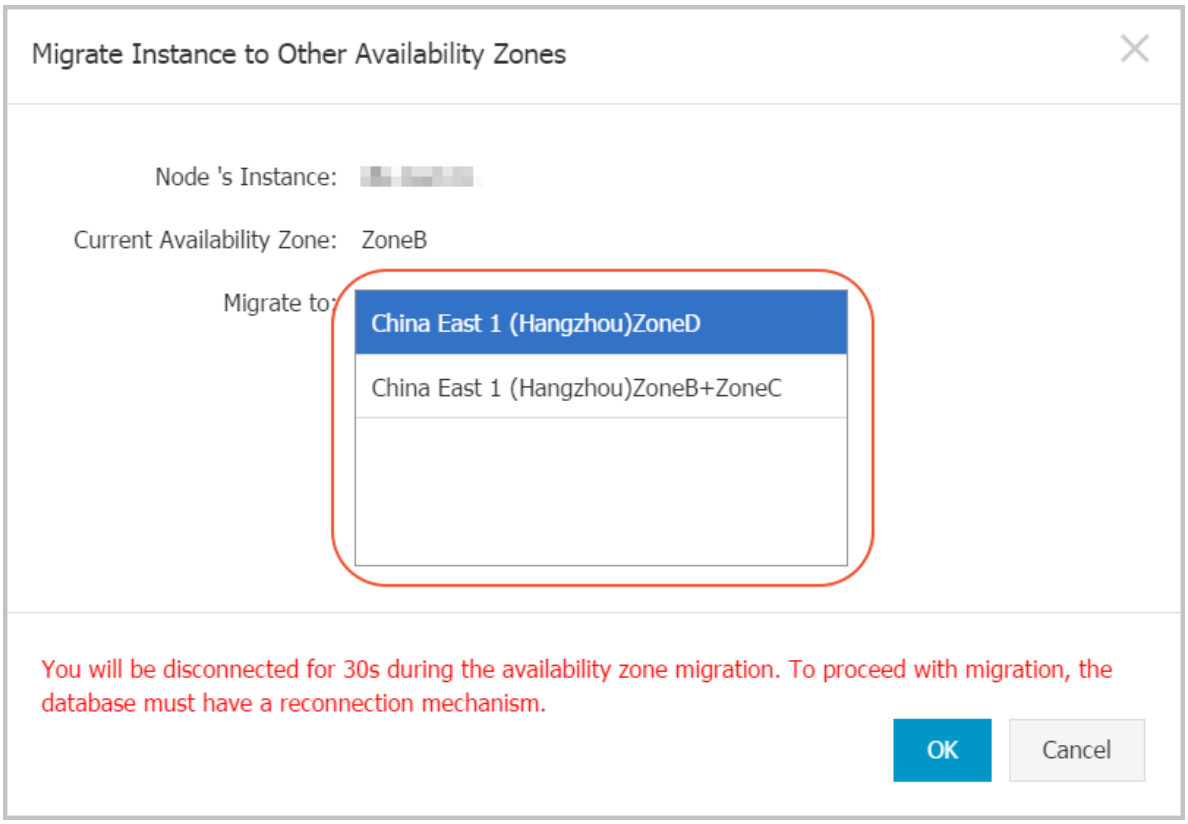

Parameter description:

- `Migrate to`: Select the region to which you want to migrate the instance.
- `Switching Time`: Choose when to perform the migration. During the migration, many operations cannot be performed. You can choose to switch immediately or at a later time.

6. To modify the maintenance time, perform the following steps. Alternatively, you can also leave the maintenance time unchanged.

   a. Click Modify, as shown in the following figure. The Basic Information page is displayed.

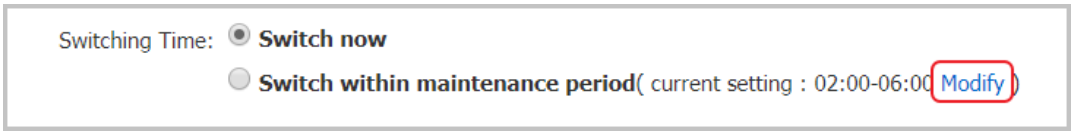

   b. At the lower left corner, select a maintenance period in the `Configurat ion Informatio n` area and click Save.

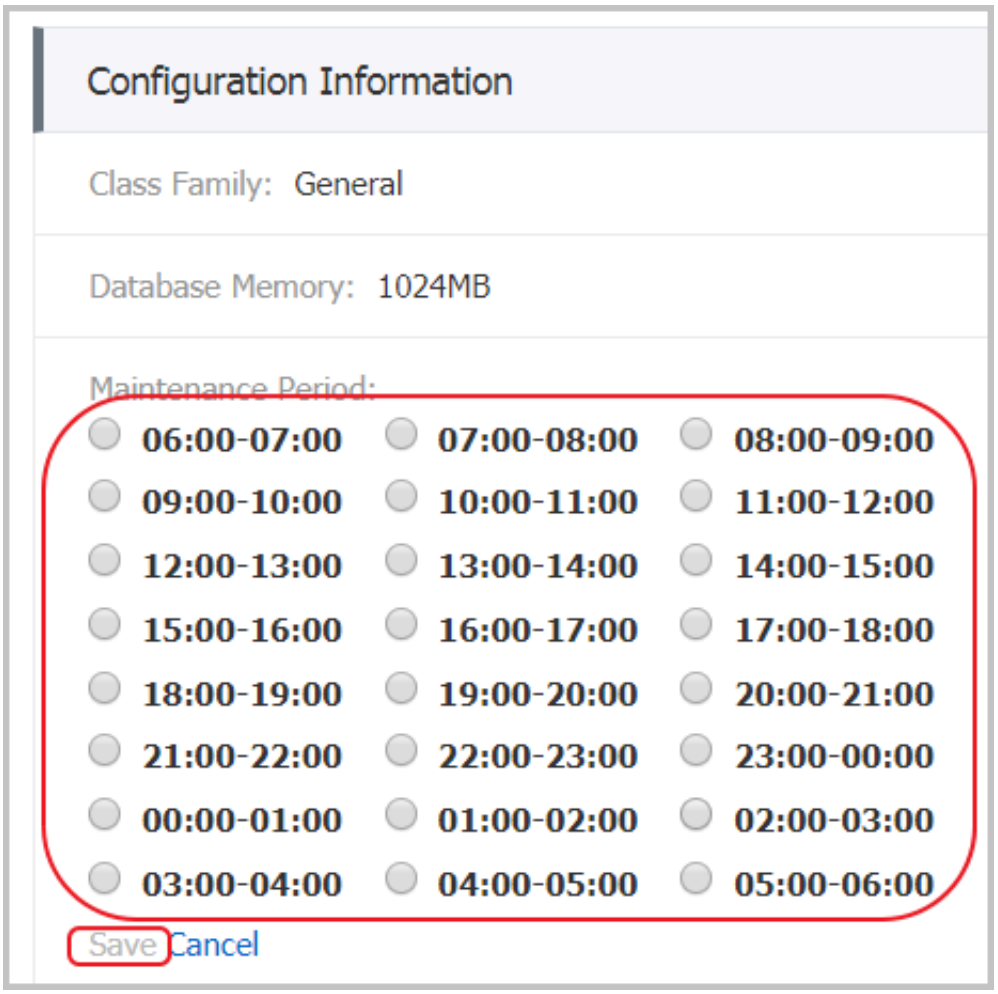

   c. Go back to the page for migrating the instance to another zone.

7. In the Migrate Instance to Other Availability Zones dialog box, click OK.

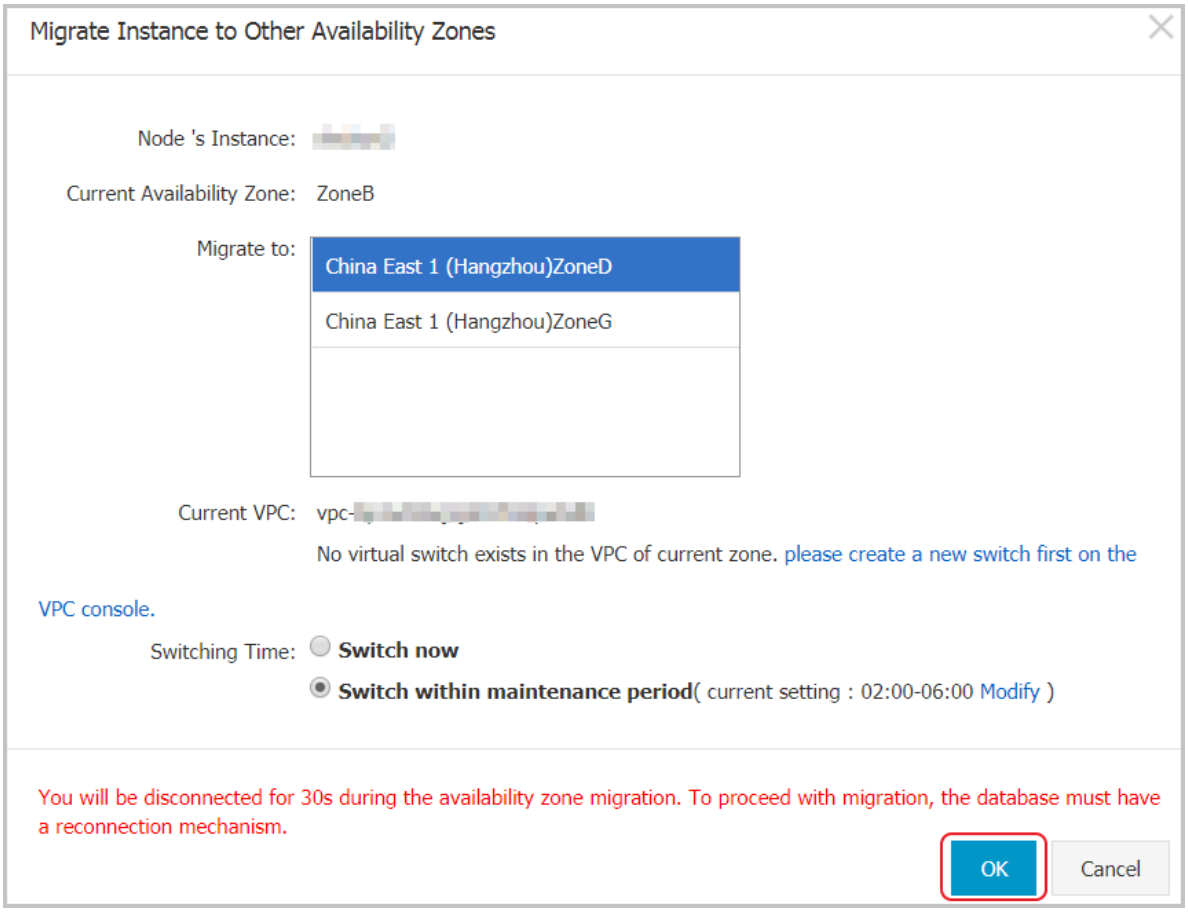

5.4 Switch between master and slave instances

Each high-availability instance consists of a master instance and a slave instance. The master and slave instances are located in different zones within the same region.

The data in the master instance is synchronized to the slave instance in real time. You can only access the master instance. The slave instance exists only as a backup. However, when the rack (where the master instance is located) encounters an error, the master and slave instances can be switched. After the switch, the original master instance becomes a backup instance, and rack-level disaster tolerance can be realized .

This document describes how to switch between master and slave instances.

Attentions

- Currently this operation is not applicable to the Basic Edition of MySQL 5.7 and SQL Server 2012/2016 instances. This is because Basic Edition instances do not have slave nodes.

- Switching between master and slave instances may result in transient disconnect ion. Make sure that your application has a reconnection configuration.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located, and click the ID of a target instance.
3. In the left-side navigation pane, select Instance Availability.
4. In the `Availabili ty Informatio n` area, click Switch Master/Slave Instance.
5. Select Switch now or Switch within maintenance period.

> **Note:**
> During the switch, many operations cannot be performed. Therefore, we recommend that you choose to switch within the maintenance period.

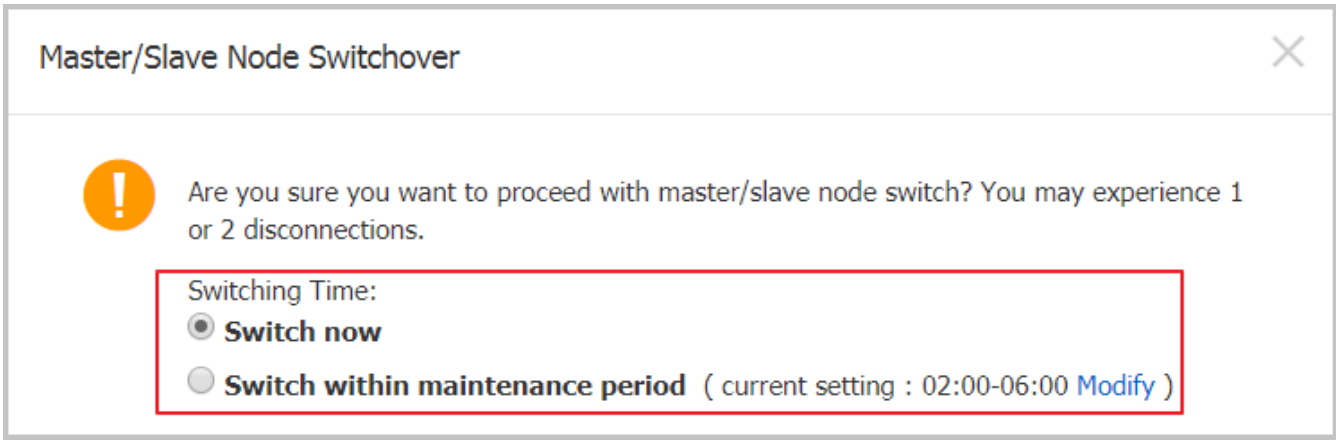

6. Do as follows to change the maintenance period is necessary:

   a. Click Modify to open the Basic Information page.

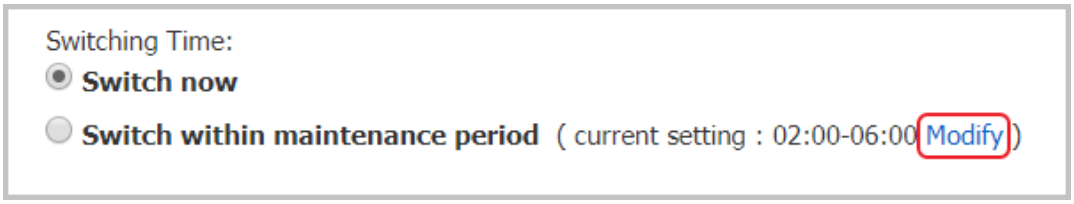

   b. In the Configuration Information area at the lower left corner, select a maintenance period and click Save.

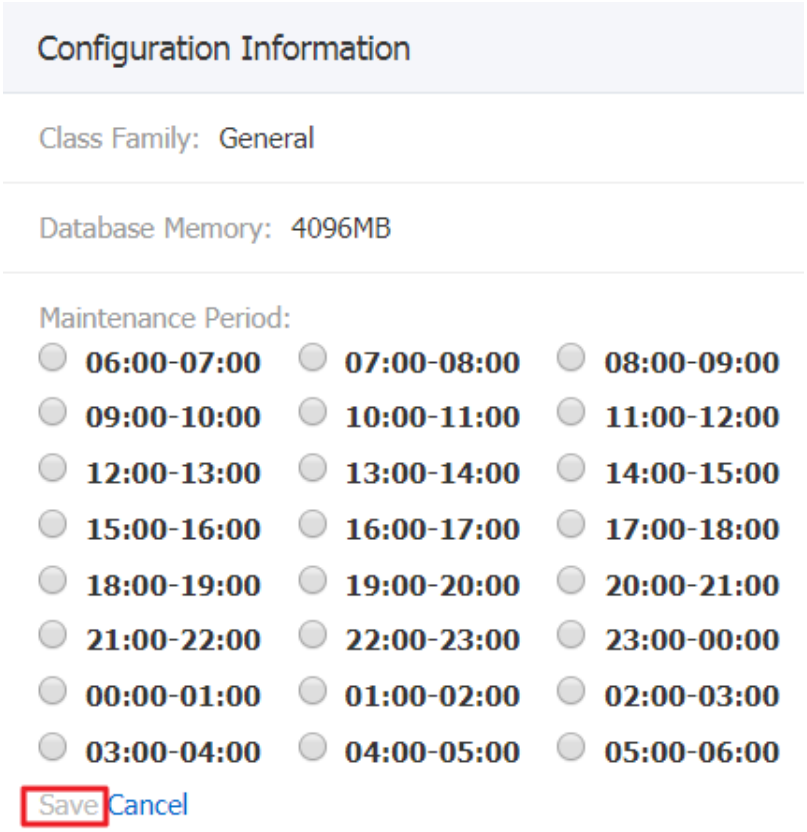

   c. Go back to the page for switching between master and slave instances and refresh the page.

7. Click OK.

5.5 Modify the data replication mode

For MySQL 5.5/5.6/5.7 instance, you can select its data replication mode based on your business characteristics to improve the availability of the RDS instance. This document introduces how to change the data replication mode.

Background information

MySQL 5.5/5.6/5.7 instances support two replication modes: semi-sync and async. You can select an appropriate replication mode as your business needs. The differences and features of the replication modes are described as follows.

- Semi-sync mode: Normally data is replicated in the sync mode. But if an exception occurs when the master node replicates data to the slave node, the data synchronization logic changes to the following:
  - When the slave node is unavailable or any network exception occurs between the master and slave nodes, the master node suspends response to applications until the replication mode times out and degrades to the async mode.
  - When data replication between the two nodes resumes normally (the slave node or network connection is recovered), async mode is changed to sync mode . The time period required for restoration to the sync mode depends on the implementation mode of the semi-sync mode. ApsaraDB for MySQL 5.5 differs from ApsaraDB for MySQL 5.6 in this regard.
- Async mode: An application initiates an update (including addition, deletion, and modification operations) request. After completing the corresponding operation, the master node immediately responds to the application and then replicates data to the slave node asynchronously. Therefore, in the async mode, unavailability of the slave node does not affect the operation on the slave database, and unavailability of the master node has a low probability to cause data inconsistency between the two nodes.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to visit the Basic Information page.
4. In the left-side navigation pane, select Instance Availability.

5. Click Modify Data Replication Mode, as shown in the following figure.

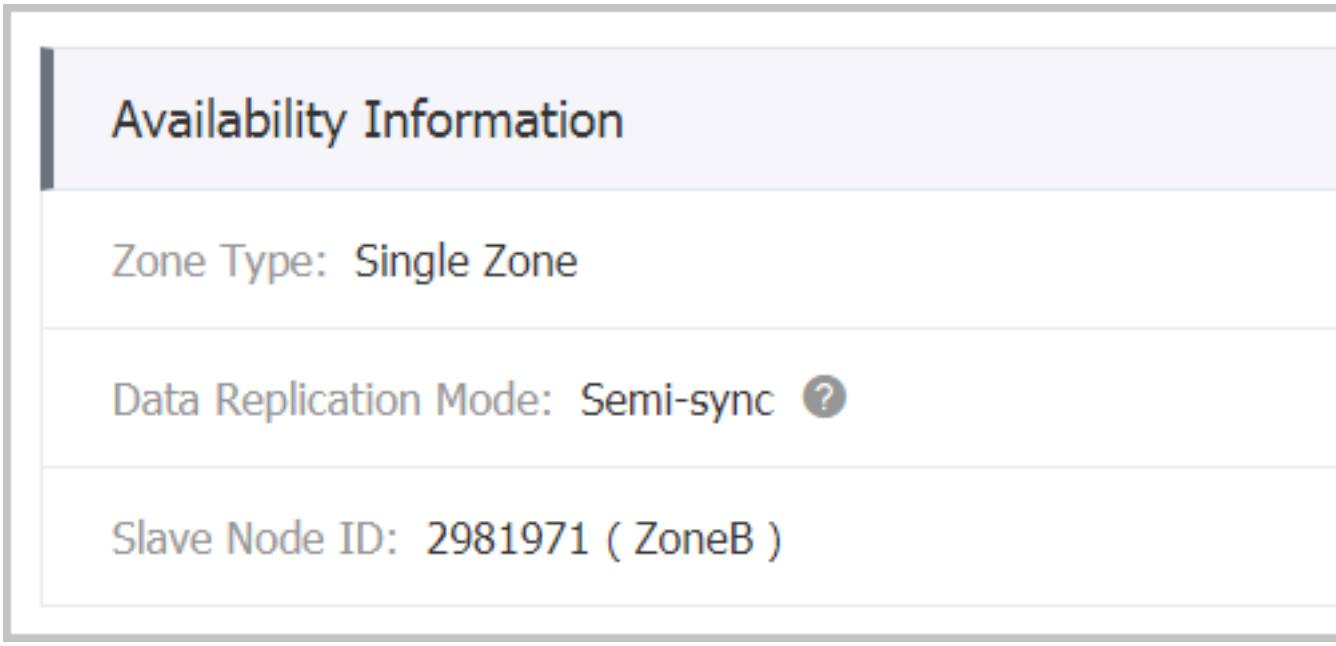

6. In the Modify Data Replication Mode dialog box, select a data replication mode, as shown in the following figure.

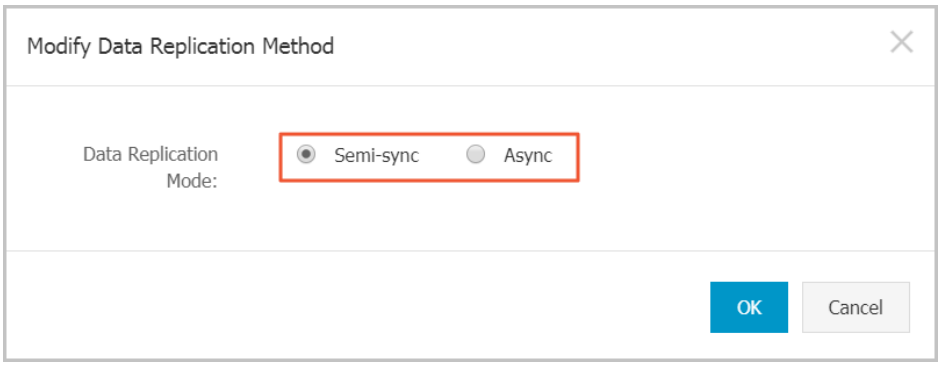

7. Click OK.

5.6 Create a read-only instance

You can create read-only instances to process massive read requests sent to the database and increase the application throughput. A read-only instance is a read-only copy of the master instance. Changes to the master instance are also automatica lly synchronized to all relevant read-only instances through the native replication capability of MySQL.

Attention

- Currently the following instances support read-only instances:
  - MySQL 5.7 High-Availability Edition (based on local SSDs)
  - MySQL 5.6
  - SQL Server 2017
- Quantity of read-only instances

Database type	Memory	Max number of read-only instances
MySQL	⩾ 64 GB	10
	< 64 GB	5
SQL Server	Any	7

- Read-only instance is subject to an additional charge and its billing method is Pay-As-You-Go. For more information, see *Pricing* for read-only instances.
- The read-only instance automatically copies the whitelist its master instance, but the whitelist of the read-only instance and that of the master instance are independent. To modify the whitelist of the read-only instance, see *Set a whitelist*.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to visit the Basic Information page.
4. In the Instance Distribution area, click Add Read-only Instance, as shown in the following figure.

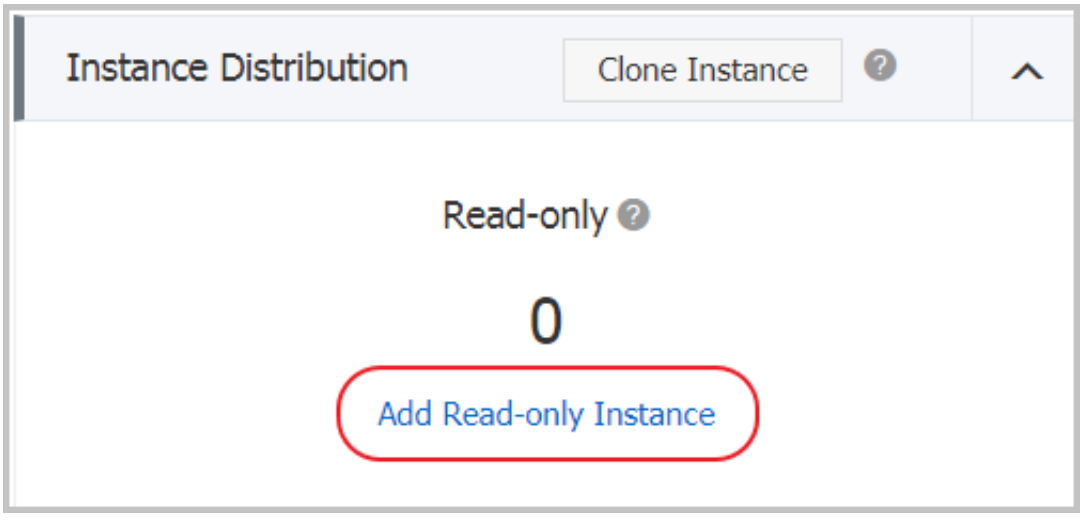

5. On the purchasing page, choose the configuration of the read-only instance, and then click Buy Now.

   Note:

   - We recommend that the read-only instance and the master instance be in the same VPC.
   - To guarantee sufficient I/O for data synchronization, we recommend that the configuration of the read-only instance (the memory) is not less than that of the master instance.
   - We recommend that you purchase multiple read-only instances to improve availability.

6. Select Product Terms of Service and Service Level Notice and Terms of Use, and then click Pay Now.
7. After creating the read-only instance, you can view it on the Instances page, as shown in the following figure.

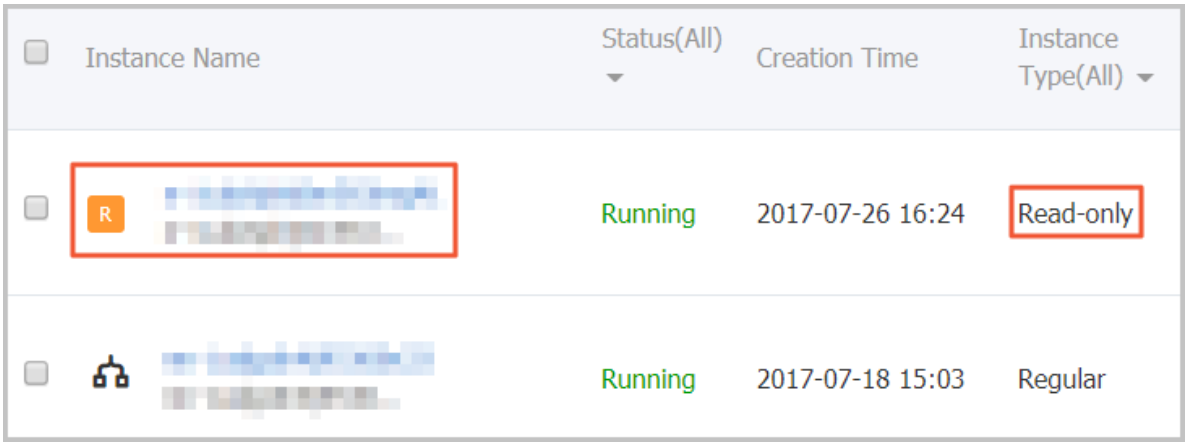

5.7 Release an instance

As your business needs change, you can manually release Pay-As-You-Go instances. This document describes detailed operations.

Attentions

- Subscription instances are released automatically when they are overdue.
- The instance is in Running status.
- For the master instance with the read/write splitting function enabled, to release read-only instances, you must *Disable read/write splitting* first.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to visit the Basic Information page.
4. In the Operating Status area, click Release Instance, as shown in the following figure.

5. In the dialog box, click Confirm to release the instance.

5.8 Upgrade the database version

Background information

RDS allows you to upgrade the database version. For more information about available target versions, see options or prompts on the RDS console.

Attentions

- Currently, this operation applies only to upgrades from MySQL 5.5 to MySQL 5.6 databases.
- We recommend that you firstly purchase an instance with the database version you want to upgrade to and verify its compatibility before upgrade.
- During the database upgrade process, the RDS service may flash off for about 30 seconds. To avoid the impacts on your production, we recommend that you upgrade the database at off-peak service hours. Alternatively, make sure that your application has the automatic reconnection policy.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to enter the Basic Information page.

4. In the Configuration Information area, click Upgrade Database, as shown in the following figure.

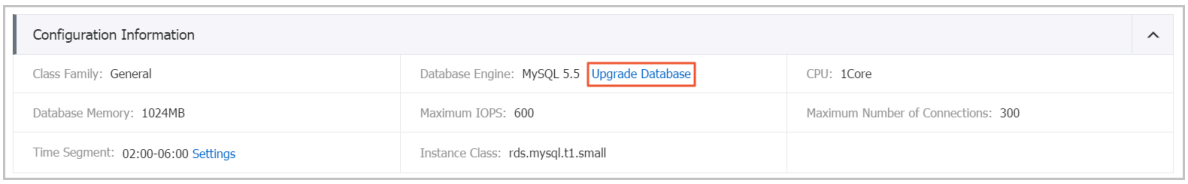

5. On the Database Version Upgrade page, select the target database version and click Start Upgrade.

5.9 RDS for MySQL release notes

MySQL 5.7

- Version 20181226
  - Supported dynamic modification of binlog-row-event-max-size to accelerate the replication of tables without primary keys.
  - Fixed the memory request problem of proxy-based instances.
- Version 20181121
  - Fixed the compatibility issue with DTS.
  - Prohibited common users to delete system databases.
- Version 20181010
  - Supported implicit primary keys.
  - Accelerated master/slave replication for tables without primary keys.
- Version 20180910

  Supported Native AIO to improve the I/O performance.
- Version 20180601
  - Prohibited non-super users from running RESET SLAVE.
  - Fixed the thread ID overflow.
- Version 20180431
  - Supported the High-availability Edition.
  - Supported *SQL explorer*.
  - Enhanced protection for instances that are generating snapshots.

MySQL 5.6

- Version 20181010

  Added the parameter rocksdb_ddl_commit_in_the_middle (MyRocks). If this parameter is turned on, certain DDL operations run the COMMIT command when being executed.

- Version 201806** (5.6.16)

  Increased the slow log precision to milliseconds.

- Version mysql_20180426 (5.6.16)

  - Supported hidden indexes so that you can set invisible indexes. For more information, see *Reference*.
  - Fixed bugs that occur when slave instances apply threads.
  - Resolved the performance deterioration that occurs when slave instances apply partition updates.
  - Resolved the problem that an entire TokuDB table is rebuilt by the ALTER TABLE COMMENT command. For more information, see *Reference*.
  - Resolved possible deadlocks triggered by the SHOW SLAVE STATUS or SHOW STATUS command.

- Version mysql_20171205 (5.6.16)

  - Resolved the problem that concurrent execution of OPTIMIZE TABLE and ONLINE ALTER TABLE causes deadlocks.
  - Resolved conflicts between SEQUENCE and implicit primary keys.
  - Resolved problems related to SHOW CREATE SEQUENCE.
  - Resolved the problem that TokuDB table statistics are incorrect.
  - Resolved the problem that parallel OPTIMIZE table commands cause deadlocks.
  - Resolved the character set problems recorded in QUERY_LOG_EVENT.
  - Resolved the problem that databases cannot be stopped due to signal processing. For more information, see *Reference*.
  - Resolved problems caused by RESET MASTER.
  - Resolved the problem that backup databases are stuck in the waiting state.
  - Resolved the possible process termination caused by SHOW CREATE TABLE.

- Version 20170927 (5.6.16)

  - Resolved the problem that TokuDB table queries use incorrect indexes.

- Version 20170901 (5.6.16)
  - Upgraded the SSL encryption version to TLS1.2. For more information, see *Reference*.
  - Supported SEQUENCE.
  - Resolved the problem that NOT IN queries return incorrect results in certain scenarios.
- Version 20170530 (5.6.16)

  Allowed master accounts to kill connections of common accounts.
- Version 20170221 (5.6.16)

  Supported *read/write splitting*.

5.10 Change configurations

As your business needs change, you can change instance configurations, that is, change instance specifications, instance series (instance changed from Basic Edition to High-availability Edition), storage space, and more. During instance configuration change:

- RDS services may experience a 30-second flash. In this case, we recommend you change instance configurations during off-peak service hours. Alternatively, make sure that your application has an automatic reconnection mechanism to avoid the impact of service burst.
- RDS allows you to set the execution time for configuration change.

Currently, only paid instances support configuration change. This document describes how to change RDS instance configuration. For information about billing of configuration changes, see *Billing details for configuration change*.

- Subscription instances:
  - During the contract period, new instance configurations (including CPU and memory) takes effect immediately after change. The number of connections and that of IOPS are increased.
  - After the instance expires, instance configurations can be upgraded or degraded during renewal. New configurations take effect at the beginning of the new billing cycle. For more information about how to renew an instance, see *Renewal*.
- Pay-As-You-Go instances can be upgraded or degraded at any time.

Attention

During configuration changes, you cannot perform most operations on databases, accounts, and networks. The following table lists the details. Choose a proper time to change instance configurations.

Function	Supported or not
Modify Whitelist	Yes
Enable SQL Audit	Yes
Set Maintenance Time Window	Yes
Add Read-only Instances	No
Add Instances for Failover	No
Release Instances	No
Switch the Billing Method to the Subscription Mode	No
Migrate Instances across Zones	No
Create User Accounts/Master Accounts	No
Reset Password	No
Change Account Permissions	No
Create and Delete Databases	No
Change Network Type	No
Change Access Mode	No
Change Connection Address	No
Apply for Internet IP Address	No
Switch between Master and Slave Instances	No
Change Backup Mode	No
Restore Data	No
Modify Parameters	No

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.

4. In the configuration information bar, click Change configuration to go to the change instance page.
5. In the change configuration bar, select a new configuration.

   Parameter description:

   - `Series` : Switch between High-availability Edition and Financial Edition instances for MySQL 5.6 and that between High-availability Edition and Basic Edition instances for MySQL 5.7 are supported.
   - `Availabili ty zone` : You can choose to migrate an instance to another availability zone, only available for MySQL 5.6 and SQL Server 2008 Release 2 instances.
   - `Specificat ions` : You can select an instance with other memory and CPU specifications.
   - `Storage` : Select the appropriate storage space based on the usage of the current database storage space.

     Note:
     The storage space of each instance specification is different, if the storage space of the current specification does not meet your needs, change the instance specifications at first, and then select the desired storage space. For more information about instance specifications, see instance spec sheets.

   - `Switch time` : Select the execution time for changing instance configurations. Changing instance configurations involves bottom-level data migration, so you can choose to change configurations immediately after the data migration is complete. There are a number of operations that cannot be performed in the event of a change, such as managing databases and accounts, switching network types. You can also change configurations during the maintenance period.

6. Do as follows if you want to modify the maintenance period. Otherwise, skip the steps.
   a. Click Modify, as shown in the following figure. The system opens a new page and turns to the Basic Information page of the instance.

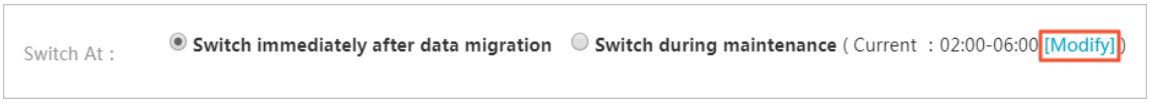

   b. In the Configuration Information area, select the maintenance period, and then click Save, as shown in the following figure.

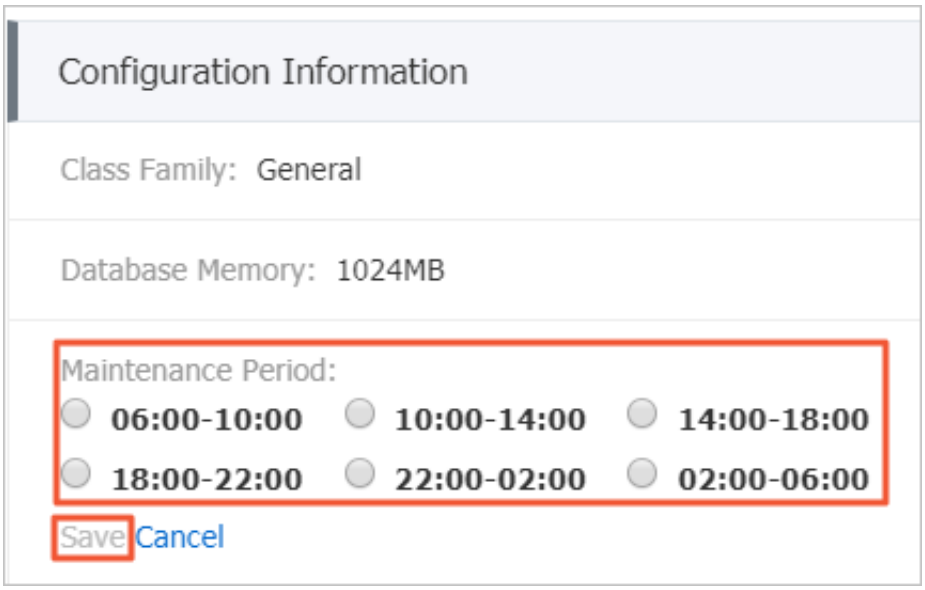

   c. Returns to the page for changing instance configurations.
7. On the instance configuration change page, click Confirm. For Subscription instances, complete the payment process according to subsequent prompts.

5.11 SQL Server DBCC function

RDS for SQL Server 2012 and later versions supports some features related to Database Console Commands (DBCC). You only need to use the stored procdure sp_rds_dbcc_trace to specify the trace flag that you want to enable. You can run `DBCC tracestatu s (- 1 )` to check whether a trace flag is enabled.

Currently, RDS supports the following trace flags:

- 1222
- 1204
- 1117

- 1118
- 1211
- 1224
- 3604

To use DBCC, run the following commands:

```
 USE   master
 GO
-- database   engine   edtion
 SELECT   SERVERPROP  ERTY (' edition ')
 GO
-- create   database
 CREATE   DATABASE   testdb
 GO

 DBCC   tracestatu  s (- 1 )

 exec   sp_rds_dbc  c_trace   1222 , 1

 WAITFOR   DELAY  ' 00 : 00 : 10 '

 DBCC   tracestatu  s (- 1 )
 GO
```

5.12 End connections for SQL Server instances

Note:

The operation described in this document is applicable only to instances of RDS for SQL Server 2012 and later versions.

Instances of RDS for SQL Server 2012 and later versions are granted the end connection (kill) permission. However, you can only end the connection that you created, for example, backup connection.

Run the following command to end a connection: `KILL ( SPID )`

6 Read/write splitting

6.1 Introduction to read/write splitting

Functions

When read/write splitting is enabled, three types of connection addresses are available in the instances:

- The connection address of the master instance: The Internet and intranet IP addresses can coexist.
- The connection address of the read-only instance: The Internet and intranet IP addresses can coexist.
- The read/write splitting address: The Internet and intranet IP addresses cannot coexist. An intranet IP address is generated by default. If you must use the read/write splitting Internet IP address, you can switch the IP address. For more information, see *Switch read/write splitting address type*.

The master instance and read-only instance require independent connection addresses. Currently, instance connection addresses are configured automatically in applications to split read and write operations.

With this function, an extra read/write splitting address is provided to associate the master instance with all its read-only instances, achieving an automatic read/write splitting link. Applications can perform read and write operations with a single connection address. Write requests are routed automatically to the master instance and read requests are routed to each read-only instance by user-defined weights. You can scale up the processing capacity of the system by adding more read-only instances without any change to applications.

The following figure shows different types of connection addresses.

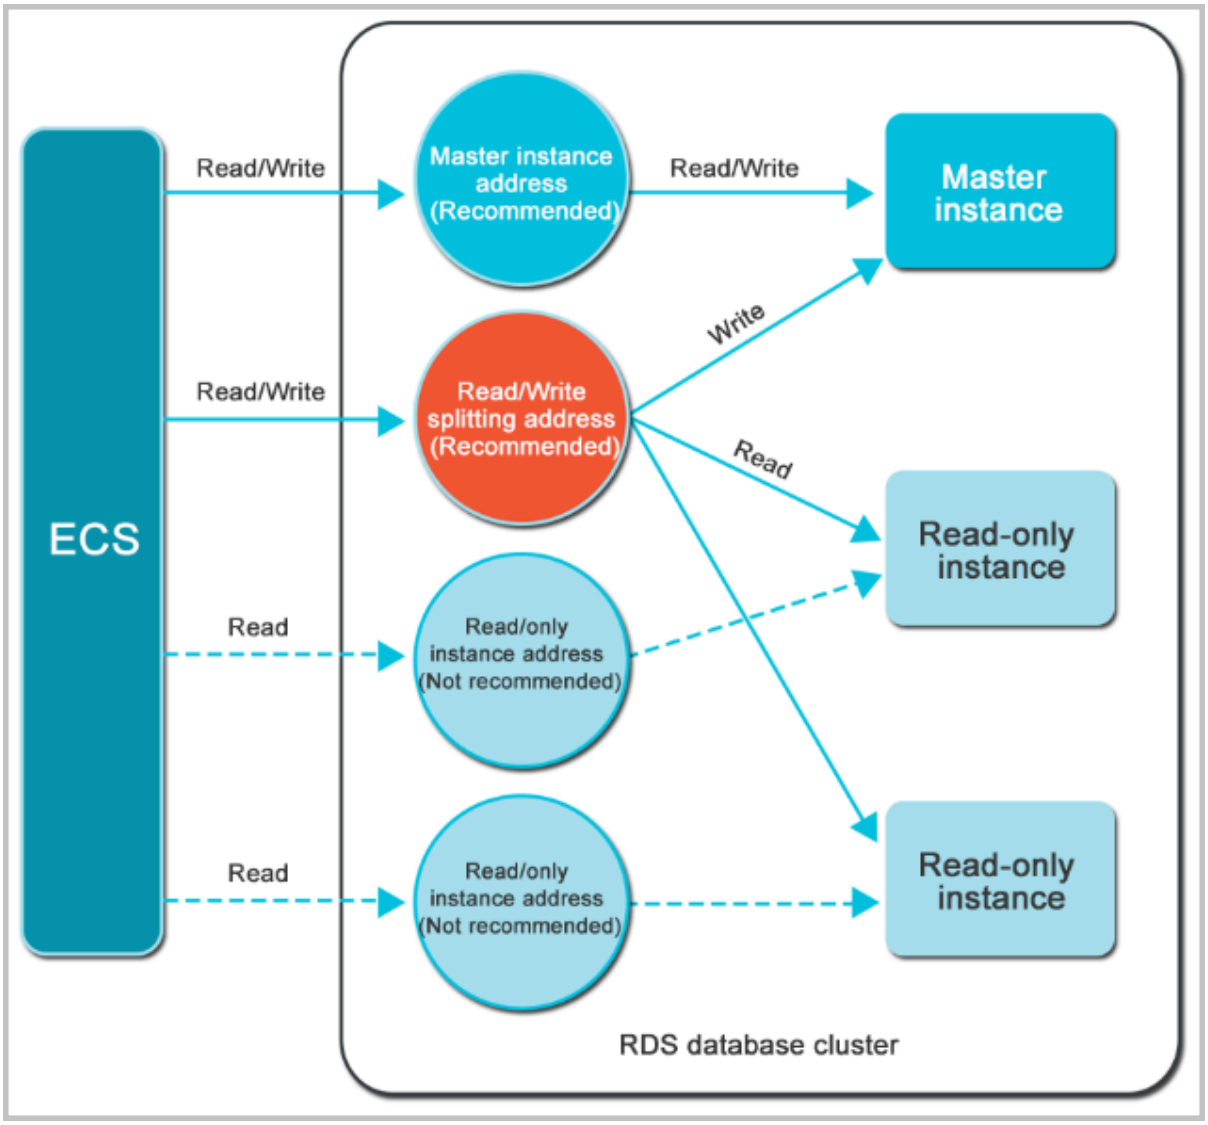

Benefits

- Facilitates maintenance with a single read/write splitting address.

  The master instance and each read-only instance have an independent connection address. You need to configure each of these addresses in applications so that write requests are sent to the master instance and read requests to read-only instances.

  The read/write splitting function provides an additional address called read/write splitting address. You can connect to this address to perform read and write operations on the master and read-only instances, with read and write requests automatically distributed. Therefore, maintenance costs are reduced.
- Improves performance with support for the highly secure link.

  For users who build a proxy layer to implement read/write splitting on the cloud, data has to go through multiple components for statement parsing and forwarding

before it reaches the database, significantly increasing the response latency. RDS read/write splitting can be directly set in the existing highly secure link without time consumption by any other components, which reduces the latency and improves the processing rate.

- Applies to various scenarios with customizable weights and thresholds.

  RDS read/write splitting can be used to set read request weights for master and read-only instances and latency thresholds for read-only instances.

- Enhances database availability with instance health checks.

  RDS read/write splitting performs health check automatically for all instances in the distribution system. If any instance fails or its latency exceeds the threshold , RDS automatically removes the instance out of the distribution system (while marking it as unavailable and stopping allocating read requests to it) and allocates read and write requests to the remaining healthy instances by the predefined weights. In this way, applications still run properly even if any single-node read-only instance fails. After the instance resumes, RDS automatically reclaims it into the request distribution system.

  Note:
  To prevent single node failures, we recommend that you create at least two read-only instances for each master instance if you are using read/write splitting.

- Reduces resource and maintenance costs with free services.

  RDS provides free read/write splitting function for all users.

Restrictions

- Currently, the following commands or functions cannot be forwarded to a read-only instance:
  - The `stmt prepare sql` command is automatically executed on the master instance.
  - The `stmt prepare command` command cannot be forwarded to a read-only instance before execution of `stmt close` .
  - The environment configuration variables `set global` , `set user` , and `set once` are automatically executed on the master instance.

- The following commands or functions are not supported currently:
  - SSL encryption
  - Compression protocols
  - com_dump_table and com_change_user protocols
  - `kill connection [ query ]`
  - `change user`
- The execution result is random for the following commands:

  The `show processlis t` , `show master status` , and `com_proces s_info` commands return results according to the instance connected during execution.
- All transactions are routed to the master database.
- Read/write splitting does not guarantee consistency of non-transactional reads. If you require such consistency, add hints to route query requests to the master database or encapsulate query requests into transactions.
- The `LAST_INSER T_ID ()` function is not supported. To use this function, add `hint : /* FORCE_MAST ER */, eg :/* FORCE_MAST ER */ SELECTLAST _INSERT_ID ();` to the request.

FAQ

How does read/write splitting ensure the timeliness of data reading?

6.2 Enable read/write splitting

In the business scenario that needs a small number of write requests but a large number of read requests to the database, you can enable the read/write splitting function to share the read pressure on the master instance. This article introduces how to enable read/write splitting.

Note:

Currently the read/write splitting function does not support the instances located in Asia Pacific NE 1 (Japan), Germany 1 (Frankfurt), Asia Pacific SE 2 (Sydney), Middle East 1 (Dubai) or Singapore.

Prerequisites

- The instance is MySQL 5.6 High-Availability Edition or Finance Edition, or MySQL 5.7 High-Availability Edition and is a master instance.
- The instance has at least one read-only instances. To create a read-only instance, see *Create read-only instance*.

Attention

- When you first enable the read/write splitting function, the system automatically upgrades the backend control system of the master instance and all the associated read-only instances to the latest version to make sure that your service works properly. Therefore, the master instance and read-only instances automatically restart once during the enabling process. The master instance is subject to a transient disconnection of up to 30 seconds, and the read-only instances cannot be accessed during the whole restart process. To avoid the influence of transient disconnection, we recommend that you enable read/write splitting during off-peak hours and make sure that automatic reconnection is available for your application.
- If you have restarted or made configuration changes to the master instance and read-only instances for which to enable read/write splitting for at least once since March 8, 2017, then the backend control system of these instances has automatically updated to the latest version. The system does not restart the instances again during the enabling process.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the target instance ID to enter the Basic Information page.
4. Select Connection Options in the left-side navigation pane to enter the Connection Options page.
5. Select Read/Write Splitting tab.
6. Click Enable now to enter the Configure Read/Write Splitting page.

Note:

If the instance was created before March 8, 2017, and it has not been restarted or its specifications have remained unchanged since March 8, 2017, the master and

read-only instances will be restarted once after you enable read/write splitting. On the displayed confirmation dialog box, click OK to enable read/write splitting.

7. Enter the setup information, as shown in the following figure.
   - Network Type: Read/write splitting address, which can be an intranet address or an Internet address. If the intranet address is selected, the intranet type of the read/write splitting address automatically matches with that of the master instance. For example, if the intranet type of the master instance is VPC (Virtual Private Cloud), then the intranet type of the read/write splitting is VPC as well.
   - Latency Threshold: This refers to the latency threshold of read-only instances with a value range of 0 to 7,200s. If the latency of a read-only instance exceeds this threshold, read requests are not forwarded to this instance regardless of its weight. Depending on the running of SQLs, latencies may occur in read-only instances. We recommend that you set the value to no less than 30s.
   - Read Weight Distribution: This refers to the read request weights of different instances. An instance with a higher weight ratio processes more read requests. For example, if a read/write splitting address is associated with one master instance and three read-only instances with a read weight of 0, 100, 200, and 200, respectively, it means that the master instance does not process read requests (write requests are automatically forwarded to the master instance), and the three read-only instances process read requests by a ratio of 1:2:2. To set the weights, you can use either of the following modes:
     - Automatic Distribution: The system automatically distributes weights for instances according to their configurations. The new read-only instances under the master instance later is automatically added to the read/write splitting link according to the set weights without manual configuration. For the read weights of instances with different specifications, see *Rules of weight distribution by system*.
     - Customized Distribution: You can customize the read request processing weights of different instances with a value range of 0 to 10,000. If you select this mode, the weight of new read-only instances added to the master instance defaults to 0, and you have to set this parameter manually.

       Note:

To obtain real-time data with certain query statements, you can forcibly forward these statements to the master instance for execution using the Hint format. For the Hint format supported by RDS read/write splitting, see "Specify whether an SQL is sent to the master instance or a read-only instance by using Hint" in the *Rules of weight distribution by system*.

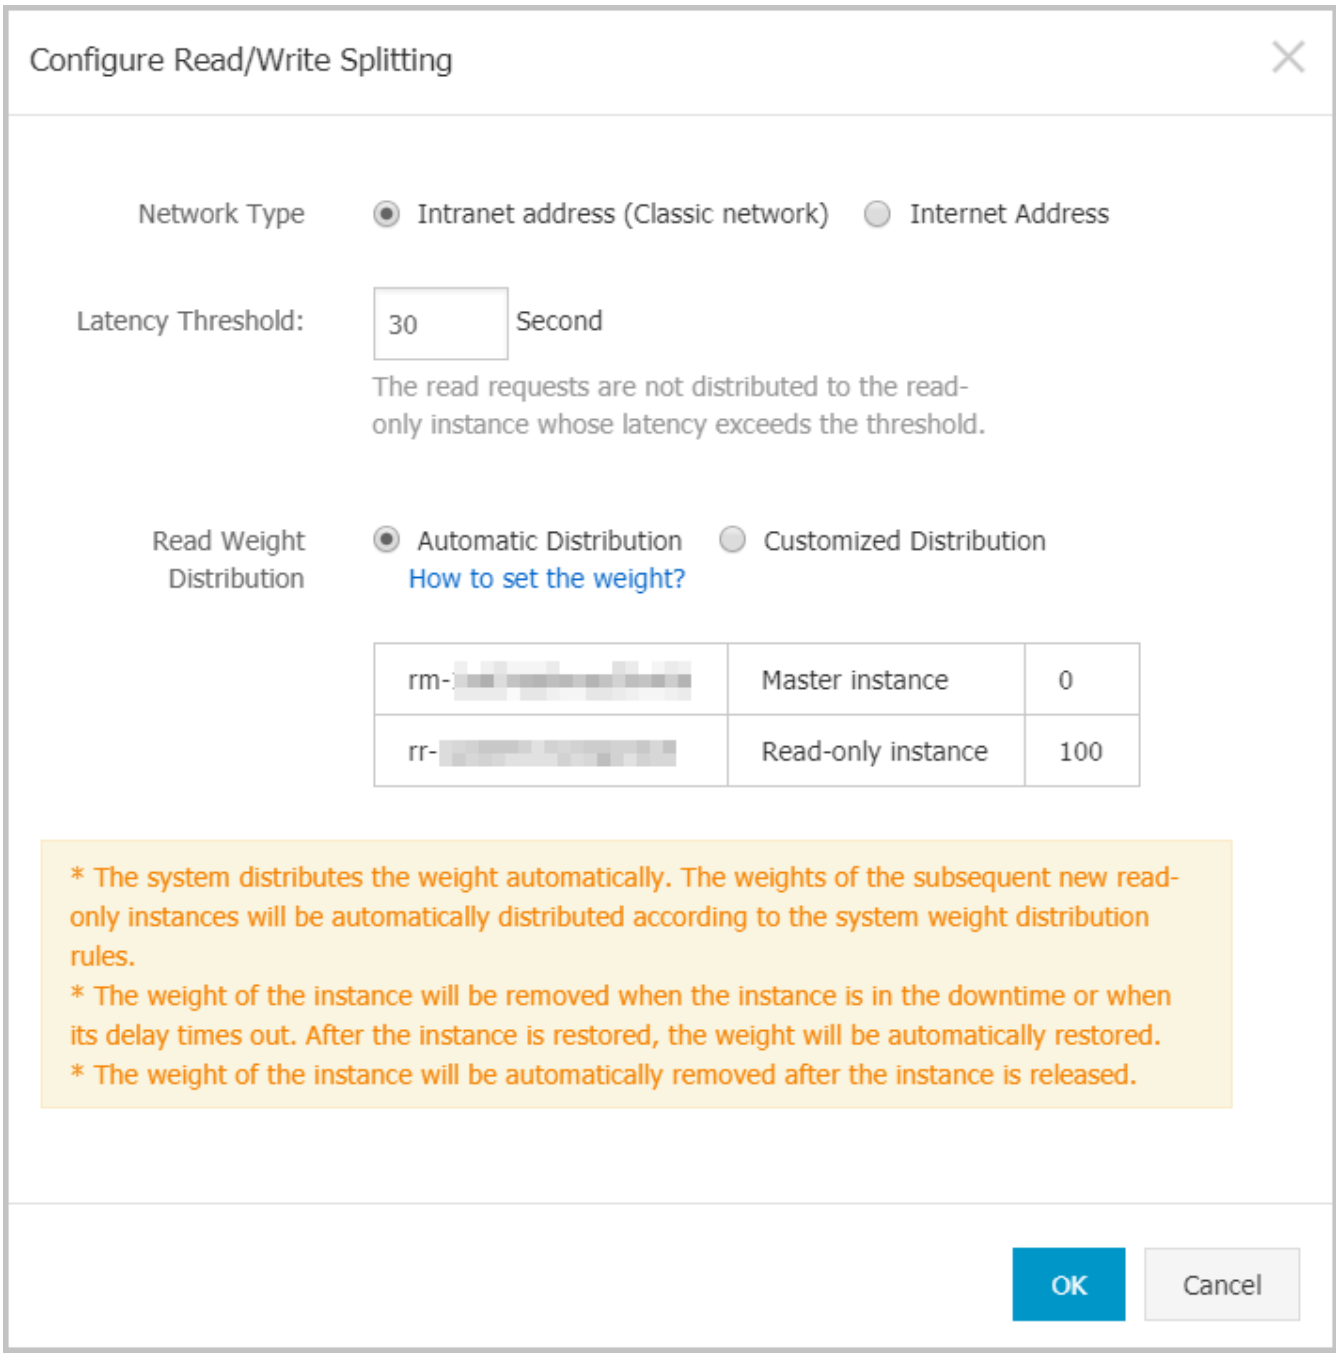

8. Click OK.

Note:

The instance status changes to Creating Network Connection. Wait for a while patiently.

6.3 Modify the latency threshold and read weights

After enabling read/write splitting, you can configure the latency threshold and read weights of instances.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to enter the Basic Information page.
4. Click Connection Options in the left-side navigation pane to enter the Connection Options page.
5. Click the Read/Write Splitting tab.

6. Click Configure Read/Write Splitting to enter the Configure Read/Write Splitting page.

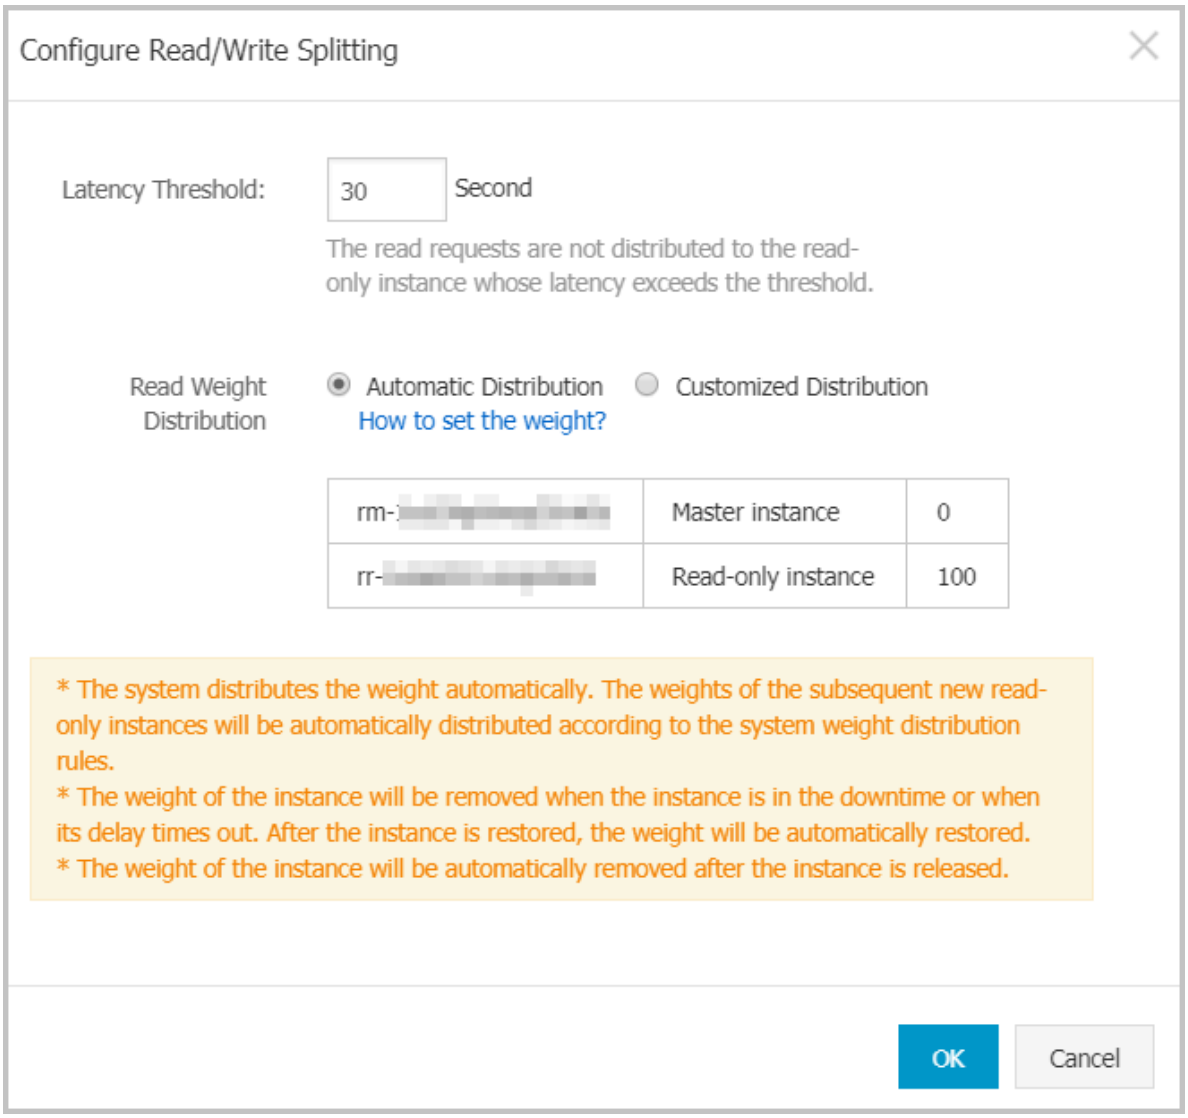

7. Change parameter settings as needed by referring to the following information:

> **Note:**
> When a read-only instance is deleted, its weight is removed automatically while the weights of other instances remain unchanged.

- `Latency Threshold` : refers to the latency threshold of read-only instances and ranges from 0-7,200s. If the latency of a read-only instance exceeds this threshold, read requests will not be forwarded to this instance regardless of its

weight. Based on SQL statement execution situations, latencies may occur in read-only instances. We recommend that you set the value to at least 30s.

- `Read Weight Distributi on` : refers to read request weights of different instances. An instance with a higher weight processes more read requests. For example, if read weights of one master instance and three read-only instances are 0, 100, 200, and 200 respectively, the master instance does not process read requests (write requests are all automatically forwarded to the master instance) while the three read-only instances process read requests with a ratio of 1:2:2. To set weights, use either of the following modes:
  - `Automatic Distributi on` : The system automatically distributes read weights to instances (including read-only instances added afterwards) according to their specifications. For more information about read weights of instances with different specifications, see *Rules of weight distribution by system*.
  - `Customized Distributi on` : You can set read weights of instances with values ranging from 0 to 10,000. In this mode, the default weight of a newly added read-only instance is 0, and you have to manually set this parameter.

> **Note:**
> To enable certain queries to return data in real time, you can forcibly forward these hint statements to the master instance. For the Hint formats supported by RDS read/write splitting, see Specify whether an SQL is sent to the master instance or a read-only instance by using Hint in *Rules of weight distribution by system*.

8. Click OK.

6.4 Switch read/write splitting address type

You can change the read/write splitting address type based on business scenarios. When you enable the read/write splitting function, the read/write splitting intranet IP address is generated by default. This document introduces how to switch between the intranet and Internet IP addresses for read/write splitting.

Prerequisite

The read/write splitting function is enabled. For more information, see *Enable read/write splitting*.

Precaution

When you change the address type, the master instance experiences a transient disconnection for up to 30 seconds. To avoid impact of the transient disconnection, we recommend that you change the network type during off-peak hours and make sure that automatic reconnection is available for your applications.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the target instance ID to enter the Basic Information page.
4. Click Connection Options in the left-side navigation pane to enter the Connection Options page.
5. Click the Read/Write Splitting tab.
6. Click Switch to Internet Address.

   Note:
   Note: If you are changing from the Internet IP address to intranet IP address, click Switch to Intranet Address.
7. In the dialog box, click Confirm.

6.5 Disable read/write splitting

If the read/write splitting function is no longer needed, you can disable it. The read /write splitting function can only be used when at least one read-only instance is available, so you must disable the read/write splitting function before you delete the last available read-only instance. Otherwise, the instance cannot be deleted.

This document explains how to disable the read/write splitting function.

Note:
After the read/write splitting function is disabled, your applications cannot connect to the read/write splitting address any longer. Make sure that your database connection configuration does not include this connection address.

Prerequisite

The instance is in MySQL 5.6 High-Availability Edition or Finance Edition, or MySQL 5 .7 High-Availability Edition with read/write splitting enabled.

Procedure

1. Log on to the *RDS console* .
2. Select the region where the target instance is located.
3. Click the target instance ID to enter the Basic Information page.
4. Click Connection Options in the left-side navigation pane to enter the Connection Options page.
5. Click the Read/Write Splitting tab.

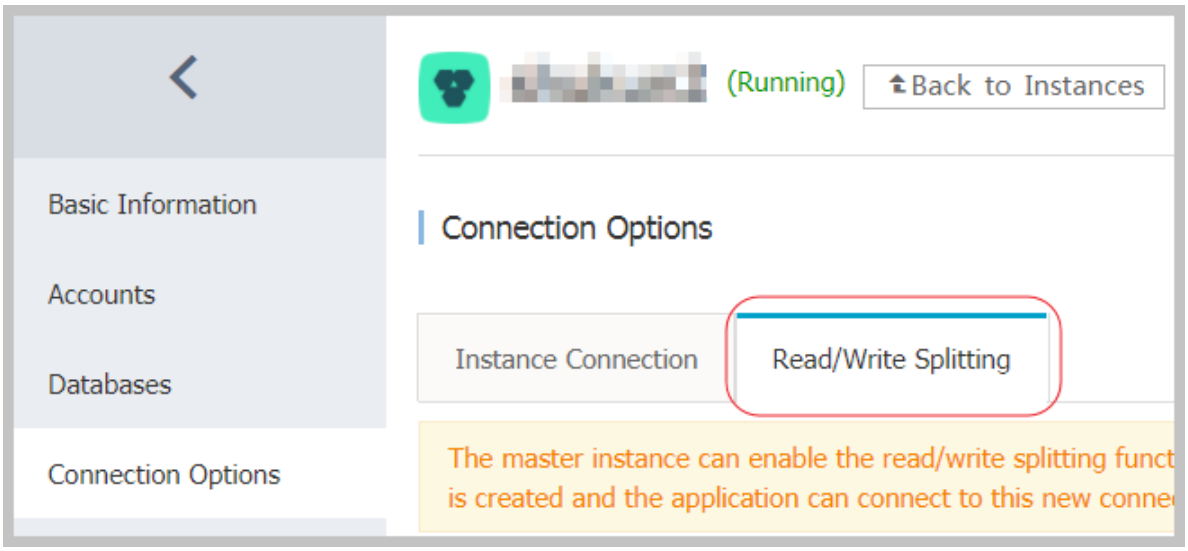

6. Click Disable Read/Write Splitting.
7. In the dialog box, click Confirm.

6.6 Monitor read/write splitting performance

You can view the read/write splitting performance on the monitoring page of the RDS console.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. Click Monitoring and Alarms in the left-side navigation pane to enter the Monitoring and Alarms page.

5. Select the monitoring type Engine Monitoring on the Monitoring tab page, as shown in the following figure.

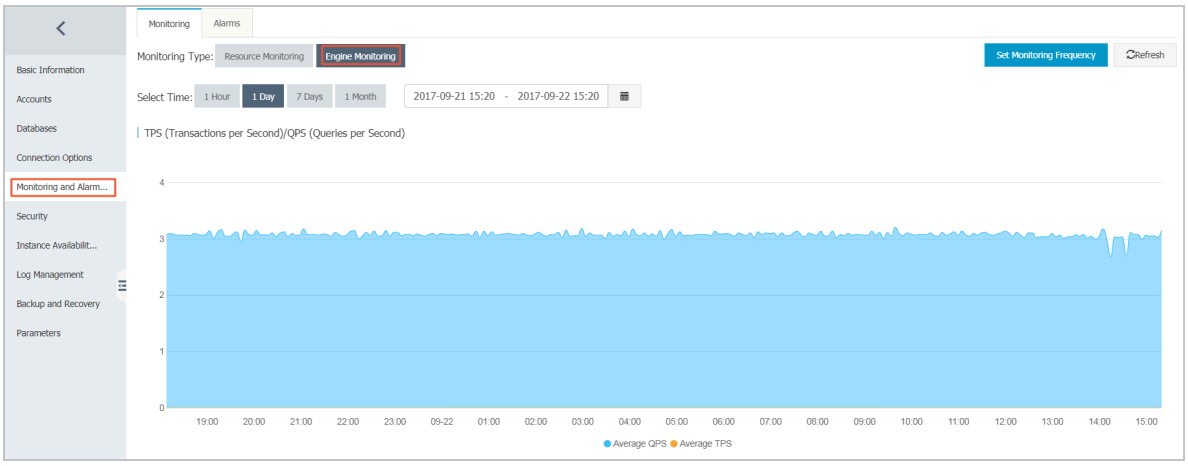

6. Query the number of read and write operations on each database (master database and read-only databases involving read/write splitting) based on Transaction Per Second (TPS) and Query Per Second (QPS).

6.7 Test read/write splitting performance

After read/write splitting is enabled, all transactions are routed to the master instance by default. Using Sysbench 0.5, the MySQL stress testing tool, as an example , this document describes how to correctly configure parameters to test read/write splitting performance.

Prerequisites

- The read/write splitting function is enabled. For detailed operations, see *Enable read/write splitting*.
- The Sysbench 0.5 is installed. Refer to *Sysbench documentation* for instructions on downloading and installing Sysbench 0.5.

Attentions

- We recommend that a case with prepare or a transaction not be for testing the load balance performance of read/write splitting.
- Prevent the master/slave latency from exceeding the threshold set for the monitoring check due to high read stress.

- We recommend that you use the following Sysbench scripts to build a specific SQL statement as needed.

```
function   thread_ini  t ( thread_id )
      db_connect ()
  end

  function   event ( thread_id )
      rs  =  db_query (" select   1 ")
  end
```

Set Sysbench parameters

A transaction is used by default to test the Sysbench oltp.lua script. If you use default parameters, all SQL statements are executed in the transaction and read-only SQL statements are routed to the master database for execution. Therefore, when the Sysbench is used to test read/write splitting performance, you must set the Sysbench parameters as needed. For example, you can set the `oltp - skip - trx` parameter to make sure that the Sysbench does not run the SQL statement in a transaction.

Set common parameters

You can set the following parameters as needed.

Name	Description
test	Path of the test file
mysql-host	IP address of the MySQL server
mysql-port	Port of the MySQL server
mysql-user	User name
mysql-password	Password
mysql-db	Database for testing, which must be created in advance
oltp-tables-count	Number of created tables
oltp-table-size	Number of records generated in each table
rand-init	Whether data is randomly initialized
max-time	Stress testing duration
max-requests	Total number of requests within a stress testing duration
num-threads	Number of concurrent threads

Name	Description
report-interval	Reporting interval of operating logs

Set parameters for transactions and read/write SQL statements

The following parameters can affect transactions and read/write SQL statements. Therefore, you must set parameters in read/write splitting tests as needed.

Name	Description
oltp-test-mode	Indicates the test mode. This parameter is unavailable in Sysbench 0.5, so this parameter can be ignored. Possible values: · `complex` : Default value. For transactional tests. · `simple` : Simple test for read-only SQL statements. · `nontrx` : For non-transactional tests. · `sp` : Stored procedures.
oltp-skip-trx	Indicates whether "begin" and "commit" of SQL statements are omitted. Possible values: · `off` : Default value. All SQL statements are executed in transactions. · `on` : Non-transactional mode. If a comparative stress test is executed repeatedly, you must run `prepare` and `cleanup` again. Note: When a stress test is executed to test the read/write splitting performance, you must set it to `on` and omit "begin" and "commit" of SQL statements.

Name	Description
oltp-read-only	Indicates whether read-only SQL statements are generated. Possible values: · `off` : Default value. The mixed read/write SQL statements of `oltp . lua` is executed. · `on` : Only read-only SQL statements are generated. UPDATE, DELETE, and INSERT SQL statements are not applicable. Note: Set the parameter value as needed to perform read-only or read/write tests.

Stress testing examples

Test read/write performance

1. Run the following command to prepare data:

```
sysbench  -- test =./ tests / db / oltp . lua  -- mysql - host =
127 . 0 . 0 . 1  -- mysql - port = 3001  -- mysql - user = abc
-- mysql - password = abc123456  -- mysql - db = testdb  -- oltp
- tables - count = 10  -- oltp - table - size = 500000  -- report
- interval = 5  -- oltp - skip - trx = on  -- oltp - read - only
= off  -- rand - init = on  -- max - requests = 0  -- max - time =
300  -- num - threads = 100  prepare ;
```

2. Run the following command to conduct the test:

Note:

When data is updated for non-transactional read/write tests, errors such as `ALERT : Error 1062 Duplicate entry ' xxx ' for key ' PRIMARY '` may occur. You must add `-- mysql - ignore - errors = 1062` to skip these errors. If the parameter `mysql - ignore - errors` does not take effect, your current Sysbench version is too old and you must upgrade it to the latest version.

```
sysbench  -- test =./ tests / db / oltp . lua  -- mysql - host =
127 . 0 . 0 . 1  -- mysql - port = 3001  -- mysql - user = abc
-- mysql - password = abc123456  -- mysql - db = testdb  -- oltp
- tables - count = 10  -- oltp - table - size = 500000  -- report
- interval = 5  -- oltp - skip - trx = on  -- oltp - read - only =
```

```
off  -- mysql - ignore - errors = 1062  -- rand - init = on  -- max
- requests = 0  -- max - time = 300  -- num - threads = 100   run ;
```

3. Run the following command to clean up data:

```
sysbench  -- test =./ tests / db / oltp . lua  -- mysql - host =
127 . 0 . 0 . 1  -- mysql - port = 3001  -- mysql - user = abc
-- mysql - password = abc123456  -- mysql - db = testdb  -- oltp
- tables - count = 10  -- oltp - table - size = 500000  -- report
- interval = 5  -- oltp - skip - trx = on  -- oltp - read - only
= off  -- rand - init = on  -- max - requests = 0  -- max - time =
300  -- num - threads = 100   cleanup ;
```

Test read-only performance

1. Run the following command to prepare data:

```
sysbench  -- test =./ tests / db / oltp . lua  -- mysql - host =
127 . 0 . 0 . 1 -- mysql - port = 3001 -- mysql - user = abc  --
mysql - password = abc123456  -- mysql - db = testdb  -- oltp -
tables - count = 10 -- oltp - table - size = 500000 -- report -
interval = 5 -- oltp - skip - trx = on  -- oltp - read - only = on
 -- rand - init = on  -- max - requests = 0 -- max - time = 300 --
num - threads = 100   prepare ;
```

2. Run the following command to conduct the test:

```
sysbench  -- test =./ tests / db / oltp . lua  -- mysql - host =
127 . 0 . 0 . 1  -- mysql - port = 3001  -- mysql - user = abc
-- mysql - password = abc123456  -- mysql - db = testdb  -- oltp -
tables - count = 10  -- oltp - table - size = 500000  -- report -
interval = 5  -- oltp - skip - trx = on  -- oltp - read - only = on
 -- rand - init = on  -- max - requests = 0  -- max - time = 300
-- num - threads = 100   run ;
```

3. Run the following command to clean up data:

```
sysbench  -- test =./ tests / db / oltp . lua  -- mysql - host =
127 . 0 . 0 . 1  -- mysql - port = 3001  -- mysql - user = abc
-- mysql - password = abc123456  -- mysql - db = testdb  -- oltp -
tables - count = 10  -- oltp - table - size = 500000  -- report -
interval = 5  -- oltp - skip - trx = on  -- oltp - read - only = on
 -- rand - init = on  -- max - requests = 0  -- max - time = 300
-- num - threads = 100   cleanup ;
```

6.8 Verify read weight distribution

To verify the load ratio of each read weight, you can run the `select @@ server_id ;` command for 10,000 times using persistent connections and collect the number of each `server_id` in the output.

Alternatively, you can verify whether the load ratio of read weight is consistent with the distributed ratio using the following methods:

Verify the load ratio based on monitoring data on the console

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the target instance ID to enter the Basic Information page.
4. Click Monitoring and Alarms in the left-side navigation pane to enter the Monitoring and Alarms page.
5. Select the monitoring type Engine Monitoring on the Monitoring tab page.
6. Query the number of read and write operations on each database (master database and read-only databases involving read/write splitting) based on Transaction Per Second (TPS) and Query Per Second (QPS).

   > **Note:**
   > Refreshing TPS/QPS performance data takes about five minutes.

7. Compare the QPS/TPS of each database to verify whether the load ratio is correct.

Verify the SQL load by directly connecting to each database

You can view the number of SQL statements executed by each instance by connecting to the master database and read-only databases involving read/write splitting.

> **Note:**
> To verify this, the connection addresses of the master database and read-only databases instead of the read/write splitting address are needed.

Run either of the following commands to verify the SQL load:

```
select * from information_schema.global_status where
VARIABLE_NAME = 'COM_SELECT';
```

```
select * from information_schema.global_status where
VARIABLE_NAME = 'COM_INSERT';
```

6.9 Rules of weight distribution by system

Weight values list

When the read weights are automatically set for instances by the system, the values of these weights are fixed, as shown in the following table:

Specification code	Specification type	Memory	CPU	Weight
rds.mys2.small	Common instances	240 MB	3	100
rds.mys2.mid	Common instances	600 MB	5	100
rds.mys2.standard	Common instances	1,200 MB	6	400
rds.mys2.large	Common instances	2,400 MB	9	400
rds.mys2.xlarge	Common instances	6,000 MB	10	800
rds.mys2.2xlarge	Common instances	12,000 MB	10	800
rds.mys2.4xlarge	Common instances	24,000 MB	12	1,000
rds.mys2.8xlarge	Common instances	48,000 MB	13	1,000
rds.mysql.t1.small	Common instances	1 GB	1	100
rds.mysql.s1.small	Common instances	2 GB	1	100
rds.mysql.s2.large	Common instances	4 GB	2	200
rds.mysql.s2.xlarge	Common instances	8 GB	2	200
rds.mysql.s3.large	Common instances	8 GB	4	400
rds.mysql.m1.medium	Common instances	16 GB	4	400
rds.mysql.c1.large	Common instances	16 GB	8	800
rds.mysql.c1.xlarge	Common instances	32 GB	8	800
rds.mysql.c2.xlarge	Common instances	64 GB	16	1,600

Specification code	Specification type	Memory	CPU	Weight
rds.mysql.c2.xlp2	Common instances	96 GB	16	1,600
rds.mysql.c2.2xlarge	Common instances	128 GB	16	1,600
mysql.x8.medium.2	Dedicated instances	16 GB	2	200
mysql.x8.large.2	Dedicated instances	32 GB	4	400
mysql.x8.xlarge.2	Dedicated instances	64 GB	8	800
mysql.x8.2xlarge.2	Dedicated instances	128 GB	16	1,600
rds.mysql.st.d13	Dedicated-host instances	220 GB	30	3,000
rds.mysql.st.h13	Dedicated-host instances	470 GB	60	6,000

Specify whether a SQL statement is sent to the master instance or a read-only instance by adding a hint

In addition to the weight distribution system of read/write splitting, hints serve as a complementary SQL syntax to specify whether a SQL statement is executed on the master instance or a read-only instance.

Hints supported by RDS read/write splitting are as follows:

- `/* FORCE_MAST ER */`: specifies that a SQL statement is executed on the master instance.
- `/* FORCE_SLAV E */`: specifies that a SQL statement is executed on a read-only instance.

For example, after a hint is prefixed to the following statement, the statement is always routed to and executed on the master instance regardless of the preset weight.

```
/* FORCE_MAST ER */ SELECT * FROM table_name ;
```

7 Account management

7.1 Reset the instance password

You can reset the password on the *RDS console* if the password for the database account is lost.

> **Note:**
> For data security, we recommend you change the password on a regular basis.

Procedure

1. Log on to the *RDS console* and select the target instance.
2. Select Accounts in the left-side navigation pane.
3. On the Account List tab page, select the account whose password you want to reset and click Reset Password.

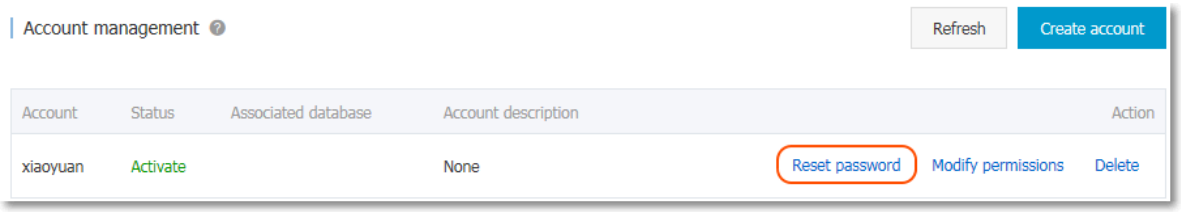

4. In the Reset Account Password dialog box, enter a new password and click OK. The password consists of 6 to 32 characters including letters, digits, hyphen (-), or underscores (_). A previously used password is not recommended.

7.2 Change account permissions

While using RDS, you can change permissions of the account at any time based on your business needs.

Procedure

1. Log on to the *RDS console* and select the target instance.
2. Select Accounts in the menu.

3. On the Account List page, find the target account and click Modify permissions, as shown in the following figure.

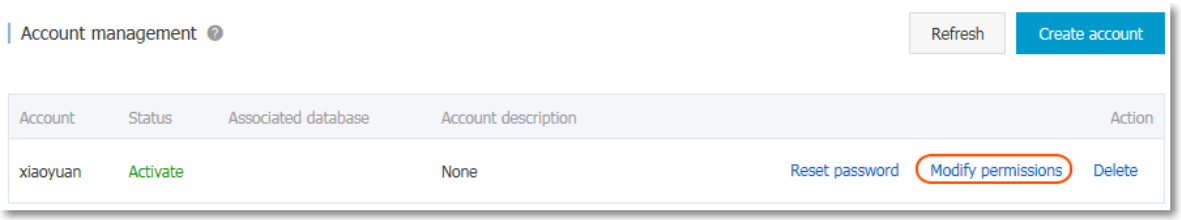

4. In the Modify account dialog box, change the account permissions and click OK, as shown in the following figure.

   - Add an authorized database: Select a database in Unauthorized database and then click Authorize > to add it to `Authorized database`.
   - Delete an authorized database: Select a database in Authorized database and then click < Removeto add it to `Unauthoriz ed database`.
   - Change permissions of Authorized database: Find a database in Authorized database and select Read/Write or Read-only. At the upper right corner of `Authorized database`, click Grant All Read/Write or Grant All Read-only.

   Note:
   Either of them is displayed at a time.

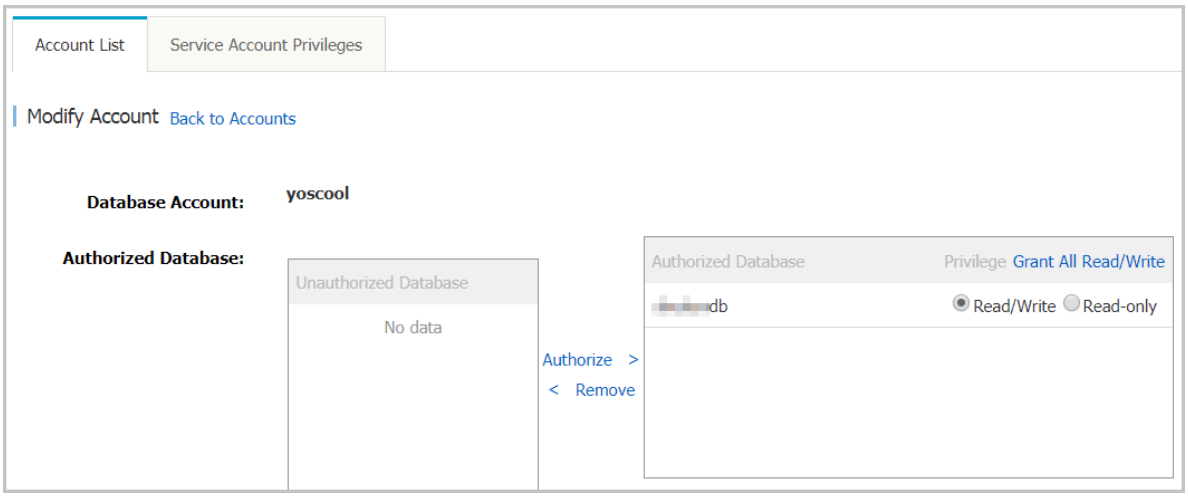

7.3 Authorize a service account

If you are seeking for technical supports from Alibaba Cloud and if it is necessary to operate your database instance during technical support, you must authorize

a service account that is used by the technical support staff to provide technical support services.

Background information

When you authorize the service account to view and modify configurations or view table structure, index, and SQL statements, the system generates a temporary service account and the corresponding permissions are given to this account according to your authorization information.

This temporary service account is automatically deleted after the validity period of authorization expires.

Procedure

1. Log on to the *RDS console* and select the target instance.
2. Select Accounts in the left-side navigation pane.
3. Select the Privilege account tab page.
4. Select the permission to be authorized to the service account and click the button in the Privilege status column, as shown in the following figure.
   - For troubleshooting of the IP whitelists, database parameters, and other problems, you must authorize Control privilege only.
   - For the database performance problems caused by your application, you must authorize Data privilege.

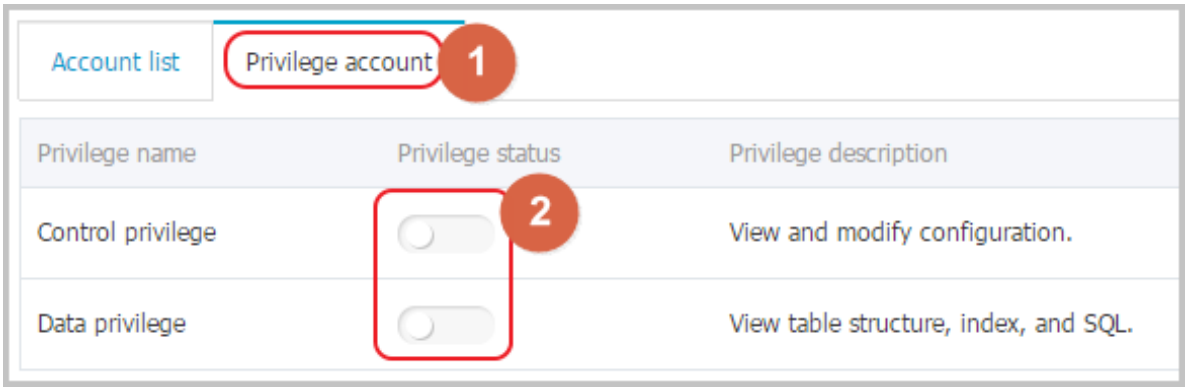

5. After setting the permission expiration time in the Setting expired time dialog box, click OK, as shown in the following figure.

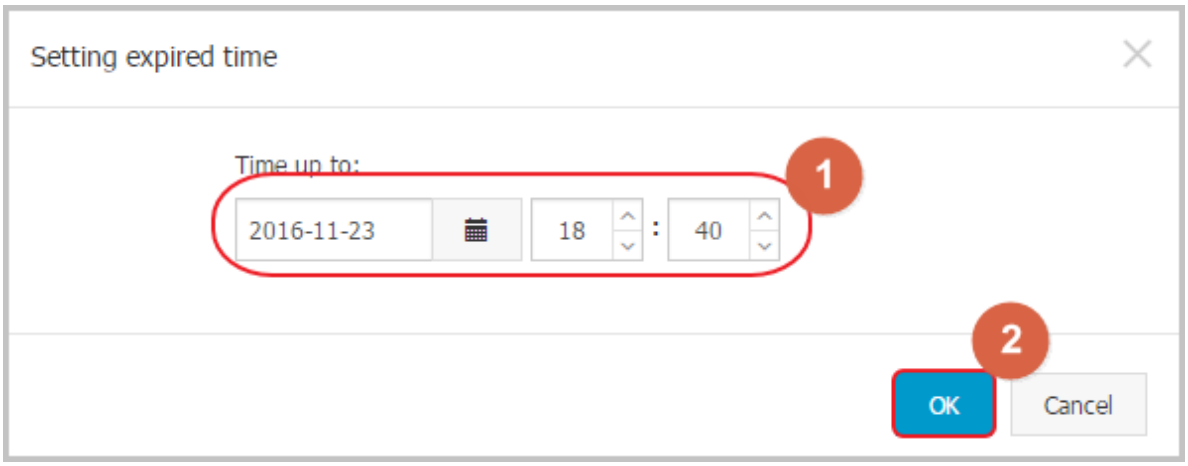

Subsequent operations

After a service account is authorized, you may cancel the authorization or change the authorization validity period on the Privilege account tab page.

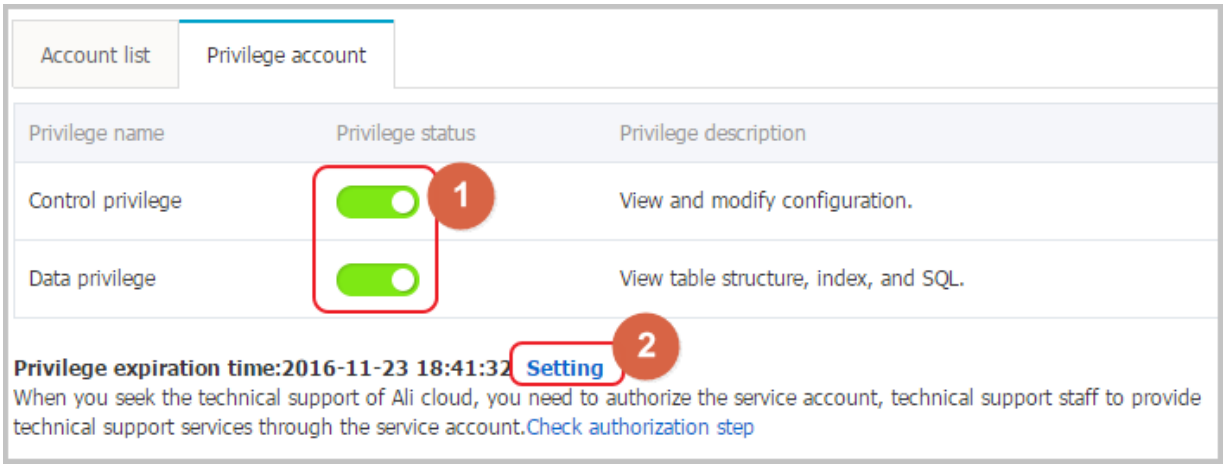

7.4 Delete an account

You can delete an account either using SQL statements or on the RDS console based on your instance type.

Delete an account on the RDS console

Currently, the RDS console allows you to delete accounts for SQL Server 2008 R2 and MySQL 5.5/5.6 instances.

Note:

If master accounts are created for MySQL 5.5 and 5.6 instances, all other common accounts can be deleted only using SQL statements.

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. In the left-side navigation pane, select Accounts to go to the Accounts page.
5. Find the account you want to delete and click Delete in its Action column.
6. In the displayed dialog box, click OK.

Delete an account using SQL statements

Currently, you can use SQL statements to delete accounts for MySQL 5.7, PostgreSQL, SQL Server 2012, and PPAS instances.

Note:
The initial or master account cannot be deleted.

1. Log on to the RDS instance. For more information, see *How to connect to ApsaraDB?*
2. Run the following command to delete the account.

```
DROP   USER  ' username '@' localhost ';
```

7.5 Manage the LOGIN user for SQL Server instances

This document describes how to create and manage the LOGIN user in a database of ApsaraDB for SQL Server.

Note:
The operation described in this document is applicable only to instances of RDS for SQL Server 2012 and later versions.

Create a LOGIN user

Run the following command to create a LOGIN user.

```
CREATE   LOGIN   Test11
WITH   PASSWORD = N ' 4C9ED138 - C8F5 - 4185 - 9E7A - 8325465CA9  B7
'
```

When the LOGIN user is being created, it is assigned permissions at the server level and database level. The Message area shows the following information.

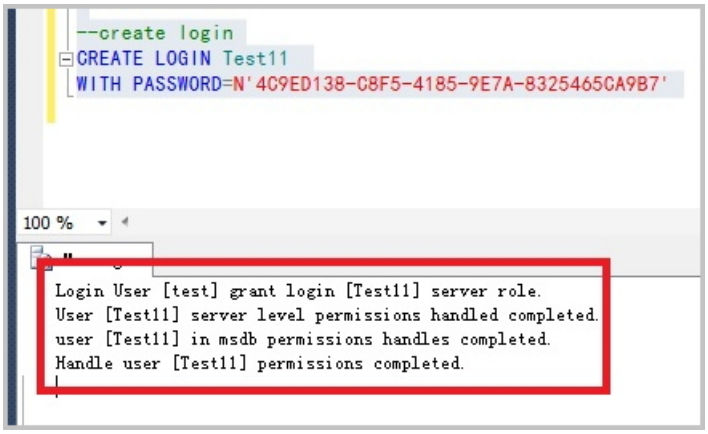

Modify a LOGIN user

Run the following commands to modify a LOGIN user.

```
ALTER   LOGIN   Test11
WITH   PASSWORD = N ' 123 ',
CHECK_POLI  CY = OFF
```

The following error is returned if you attempt to modify a LOGIN user that is not created by you.

Delete a LOGIN user

Run the following command to delete a LOGIN user:

```
DROP   LOGIN   Test11
```

An error is returned if you attempt to delete a LOGIN user that is not created by you.

7.6 Manage users for SQL Server instances

You can create common users in the database that you created other than the system database. This document describes how to create and manage users in a database of ApsaraDB for SQL Server using SQL commands.

Note:
The operation described in this document is applicable only to instances of RDS for SQL Server 2012 and later versions.

Prerequisites

- You have created a user database. For information about the commands used to create a database, see *Database management of SQL Server instances*.
- You have created a LOGIN user and logged on to the database where you plan to create a common user. For information about the commands used to create a LOGIN user, see *LOGIN user management of SQL Server instances*.

Create a user

Run the following commands to create a user in the database named TestDB:

```
USE   TestDB
Go
CREATE   USER  [ Test ]  FOR   LOGIN  [ Test ]
```

Modify user information

Modify user information in accordance with the corresponding operation instructio ns of SQL Server. For example, you can run the following commands to modify user-mapped logon information:

```
USE   TestDB
GO
ALTER   USER   test   WITH   LOGIN = test
```

Delete a user

Run the following commands to delete a user (the operation is the same as that on SQL Server):

```
USE   TestDB
GO
DROP   USER   test
```

8 Database management

8.1 Create a database

- *MySQL*
- *SQL Server 2008 R2*
- *SQL Server 2012/2016*
- *SQL Server 2017*
- *PostgreSQL*
- *PPAS*
- *MariaDB*

8.2 Manage databases of SQL Server instances

This document describes how to create and manage databases in an instance of ApsaraDB for SQL Server using SQL statements.

> **Note:**
> The operation described in this document is applicable only to instances of RDS for SQL Server 2012 and later versions.

Create a database

Run the following command to create a database:

> **Note:**
> A default path is generated when you create a database in RDS. Therefore, do not specify any file path.

```
CREATE DATABASE TestDb
```

Modify a database

You can modify many database attributes as needed. However, do not perform the following operations unless necessary:

- **Do not move the database to an incorrect file path.**

  For example, if you specify an incorrect file path by running the following commands:

```
ALTER   DATABASE  [ TestDb ]
MODIFY   FILE
(  NAME  =  N ' TestDb ',  FILENAME  =  N ' E :\ KKKK \ DDD \ DATA \
TestDb . mdf ' )
```

  The following error message will be returned:

```
Msg   50000 ,  Level   16 ,  State   1 ,  Procedure  ******,  Line
  152
The   file   path  [
E :\ KKKK \ DDD \ DATA \ TestDb . mdf  ]  is   invalid , please
specify   correct   path   folder  [  E :\ mmm \ gggg \ ].
Msg   3609 ,  Level   16 ,  State   2 ,  Line   2
The   transactio  n   ended   in   the   trigger .  The   batch
has   been   aborted .
```

- **Do not set the database recovery mode to a mode other than FULL.**

  For example, if you set the database recovery mode to SIMPLE by running the following commands:

```
ALTER   DATABASE  [ TestDb ]
SET   RECOVERY   SIMPLE
```

  The following error message will be returned:

```
Msg   50000 ,  Level   16 ,  State   1 ,  Procedure  ******,  Line
  46
Login   User  [ Test11 ]  can ' t   change   database  [ TestDb ]
recovery   model .
Msg   3609 ,  Level   16 ,  State   2 ,  Line   2
The   transactio  n   ended   in   the   trigger .  The   batch
has   been   aborted .
```

- **Do not set a database in offline state to online directly.**

  For example, if you directly run the following commands:

```
USE  [ master ]
  GO

 -- set   offline
 -- ALTER   DATABASE  [ TestDb ]
 -- SET   OFFLINE
 -- WITH   ROLLBACK   AFTER   0

  ALTER   DATABASE  [ TestDb ]
```

```
   SET   ONLINE
```

The following error message will be returned:

```
Msg   5011 ,  Level   14 ,  State   9 ,  Line   1
User   does   not   have   permission   to   alter   database
' TestDb ',  the   database   does   not   exist ,  or   the
database   is   not   in   a   state   that   allows   access
checks .
Msg   5069 ,  Level   16 ,  State   1 ,  Line   1
ALTER   DATABASE   statement   failed .
```

To change the database status from offline to online, run the following command in the `sp_rds_set _db_online` stored procedure:

```
EXEC   sp_rds_set  _db_online  ' TestDb '
```

Delete a database

Run the following command to delete a database:

```
DROP   DATABASE  [ TestDb ]
```

The following prompt appears if the database to be deleted is not backed up:

```
DROP   DATABASE  [ TestDb ]

        Kindly   reminder :
            your   database  [ TestDb ]  does   not   exist   any
backup   set .

Login   User  [ Test11 ]  has   dropped   database  [ TestDb ] .
```

9 Connection management

9.1 Set the access mode

RDS supports two access modes: standard mode and high-security mode. The high-security mode is also called database proxy mode, which may cause service instability. If your instance (excluding SQL Server 2008 R2) is in high-security mode, you need to switch to the standard mode. For more information, see *[Important] RDS network link upgrade*.

Differences between standard mode and high-security mode

- Standard mode (recommended): RDS uses *Sever Load Balancer* to eliminate the impact of database failovers on the application layer. This shortens response time and increases performance.
- High-security mode (database proxy mode): This mode prevents 90% of disconnections, but increases response time by 20% or more. If your instance (excluding SQL Server 2008 R2) is in this mode, you need to switch to the standard mode.

How to switch the access mode

Note:

- You can only disable the database proxy mode (that is, switch from the high-security mode to the standard mode), and cannot enable the database proxy mode (that is, switch from the standard mode to the high-security mode).
- During the access mode switching, the RDS instance may be disconnected once for about 30 seconds. It is recommended that you perform the switching during off-peak hours and make sure that your application can automatically reconnect to the RDS instance.

Method 1

1. Log on to the *RDS console*.
2. In the upper-left corner, select the region.
3. Locate the target instance and click the instance ID.
4. In the left-side navigation pane, click Connection Options.

5. Click Switch Access Mode.

> Note:
> This button is available only if you have turned on the database proxy mode.

Method 2

1. Log on to the *RDS console*.
2. In the upper-left corner, select the region.
3. Locate the target instance and click the instance ID.
4. In the left-side navigation pane, click Database Proxy.
5. In the Database Proxy tab page, click the slider to turn on or off the database proxy mode.

> Note:
> This tab page is available only if you have turned on the database proxy mode.

9.2 Set network type

RDS supports two network types: classic network and Virtual Private Cloud (VPC). We recommend VPC because it provides higher security. This document describes the differences between the two network types and the method of switching between the network types.

> Note:
> To migrate an instance from a classic network to a VPC without service interruptions, see *Hybrid access solution for smooth migration from classic networks to VPCs*.

Background information

On the Alibaba Cloud platform, a classic network and a VPC differs in the following aspects:

- Classic network: Cloud services in a classic network are not isolated, and unauthorized access can be blocked only by the security group or whitelist policy of cloud services.
- VPC: It helps you build an isolated network environment in Alibaba Cloud. You can customize the routing table, IP address range and gateway on the VPC. In addition, you can combine your data center and cloud resources in the Alibaba Cloud

VPC into a virtual data center through a leased line or VPN to smoothly migrate applications to the cloud.

Precautions

- After switching the network type, the original intranet IP address is changed and the Internet IP address remains unchanged. Update the connection address on your applications if necessary. For example, after an RDS instance is switched from a classic network to a VPC, the intranet IP address of the classic network is released and a VPC IP address is generated. Therefore, ECS instances in classic networks cannot access the RDS instance through the intranet any more.
- To switch MySQL 5.5, MySQL 5.6, or SQL Server 2008 R2 instances from a classic network to a VPC, the access mode must be set to safe connection mode. To switch the access mode, see *Disable database proxy Mode*.

  Note:
  MySQL 5.5, MySQL 5.6, and SQL Server 2008 R2 instances in North China 1, North China 2, East China 1, and Hong Kong regions do not have this constraint.
- During network type switching, RDS services may be interrupted for about 30 seconds. Therefore, switch the network type during off-peak hours or make sure that your applications have the automatic reconnection mechanism.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to enter the Basic Information page.
4. Click Connection Options in the left-side navigation pane to open the Connection Options page.
5. Do as follows to switch the network type:
   - Switch from a classic network to a VPC
     a. Click Switch to VPC.
     b. Select a VPC and a virtual switch.

     Note:
     - If the drop-down lists do not display VPCs or virtual switches or if the VPCs and virtual switches are not what you need, create a VPC and virtual

switch that are in the same region as the RDS instance. To create a VPC, see *Create a VPC*. To create a virtual switch, see *Create a switch*.

- For MySQL 5.5, MySQL 5.6, and SQL Server 2008 instances, their access mode must be safe connection mode if you want to switch from a classic network to a VPC. To switch the access mode, see *Disable database proxy Mode*.

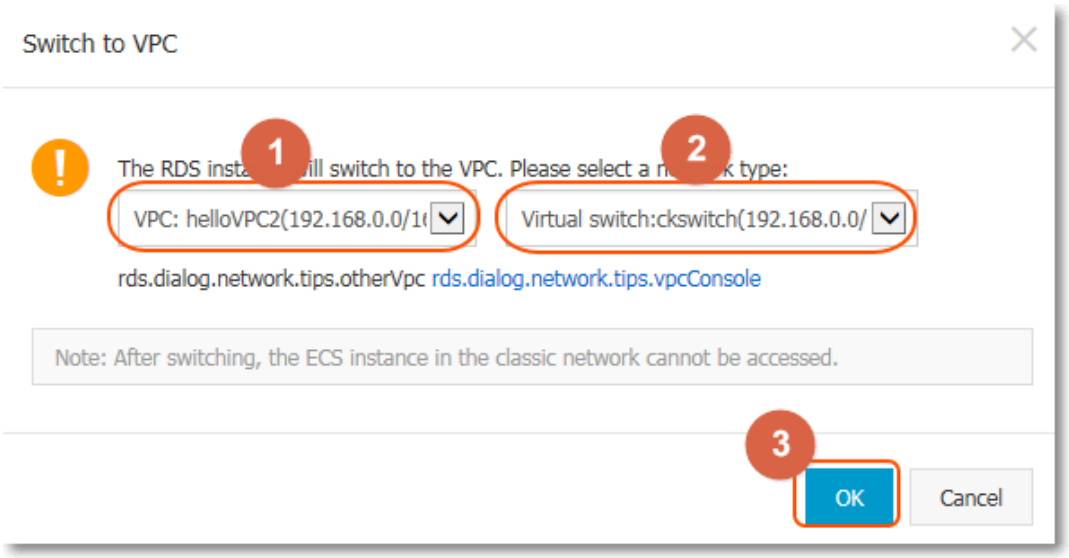

c. Click OK.

- Switch from a VPC to a classic network

a. Click Switch to Classic Network.

b. Click OK.

9.3 Hybrid access solution for smooth migration from classic networks to VPCs

Virtual Private Cloud (VPC) is a private network logically isolated from other virtual networks. A VPC allows you to build an isolated network environment with better security and performance than classic networks. With these benefits, VPCs have become a preferred networking choice for cloud users.

To meet the increasing network migration needs, RDS has added a new feature called hybrid access mode. This feature enables smooth migration from classic networks to VPCs with no intermittent service interruption or access interruption. The feature also offers the option to migrate a master instance and its read-only instances separately to a VPC without any interference with each other.

This document explains how to migrate from a classic network to a VPC on the RDS console using the hybrid access solution.

Background information

With a traditional solution, migrating an RDS instance from a classic network to a VPC causes immediate release of classic network IP address. As a result, an intermitte nt interruption for up to 30 seconds may be caused, and ECS on the classic network can no longer access the RDS instance using the intranet IP address, which may have negative impact on your services. In many large companies, a database is usually designed for access by more than one application system. When they decide to migrate the database from a classic network to a VPC, it would be quite difficult to migrate the network of all the applications simultaneously, which may result in bigger impact on their services. Therefore, a transitional period is required. To accommodate the need for smooth migration, RDS has added the hybrid access feature, making it possible to have such a transitional period.

Hybrid access refers to the ability of an RDS instance to be accessed by ECSs on both a classic network and a VPC. During the hybrid access period, the RDS instance reserves the intranet IP address of the original classic network and adds an intranet IP address for a VPC, which prevents any intermittent interruption during migration . We recommend that you use a VPC only for purposes of security and performance . For this reason, hybrid access is available for a limited period of time. That means the intranet IP address of the original classic network is released when the hybrid access period expires. In this case, your applications cannot access the database using the intranet IP address of the classic network. You must configure the intranet IP address for a VPC in all your applications during the hybrid access period to guarantee smooth network migration and minimize the impact on your services.

For example, a company wants to migrate its database from a classic network to a VPC . The hybrid access solution can be used to provide a transitional period during which some of their applications can access the database through a VPC, and the others can continue to access the database through original classic network. When all the applications can access the database through the VPC, the intranet IP address of the original classic network can be released, as shown in the following figure.

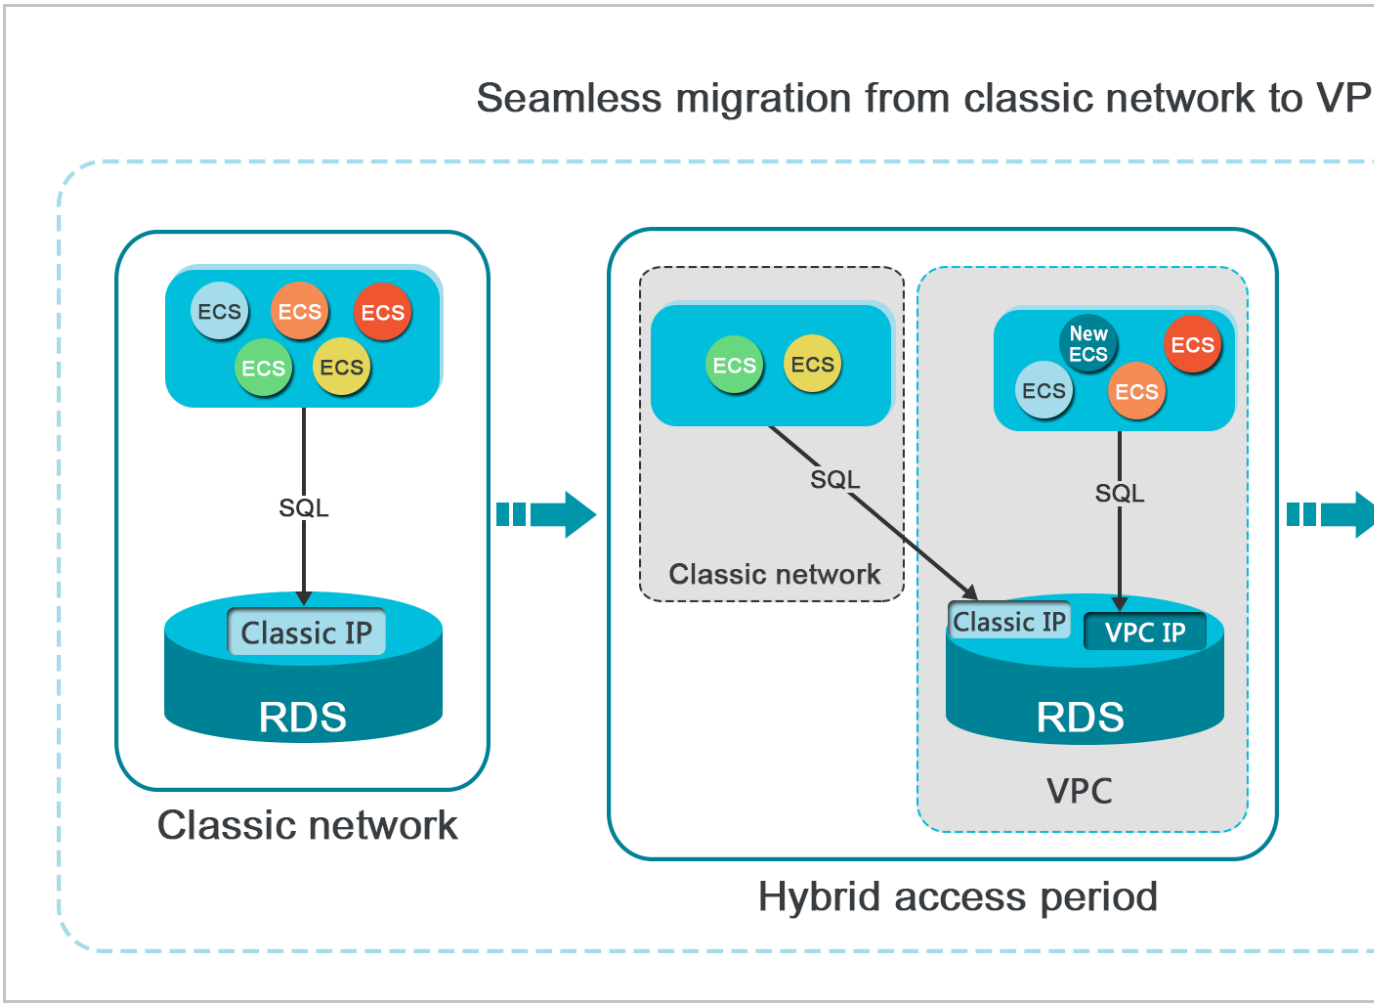

Functional Limits

The following functional limits are proposed during the hybrid access period:

- Switch to classic networks is not supported.
- Zone migration is not supported.

Prerequisites

- The current network type is classic network.
- There are available VPC and vSwitch in the zone where the RDS instance is located. If not, create them by referring to *Create VPC* and *Create vSwitch*.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.

4. In the left-side navigation pane, click Connection Options to enter the Connection Options page.
5. On the Instance Connection tab page, click Switch to VPC.
6. On the Switch to VPC confirmation page, select the target VPC and Vswitch.
7. Check Reserve original classic endpoint, and select the Expiration time for the basic intranet IP address of the original network, as shown in the following figure.

Note:

- From the seventh day before the date on which the intranet IP address of the original classic network is to be released, the system sends a text message of a notice to the mobile number bound to your account every day.
- When the reservation ages out, the intranet IP address of the classic network is automatically released and can no longer be used to access the database. To

prevent service interruption, set a reservation period as necessary. After the hybrid access configuration is complete, you can change the expiration date.

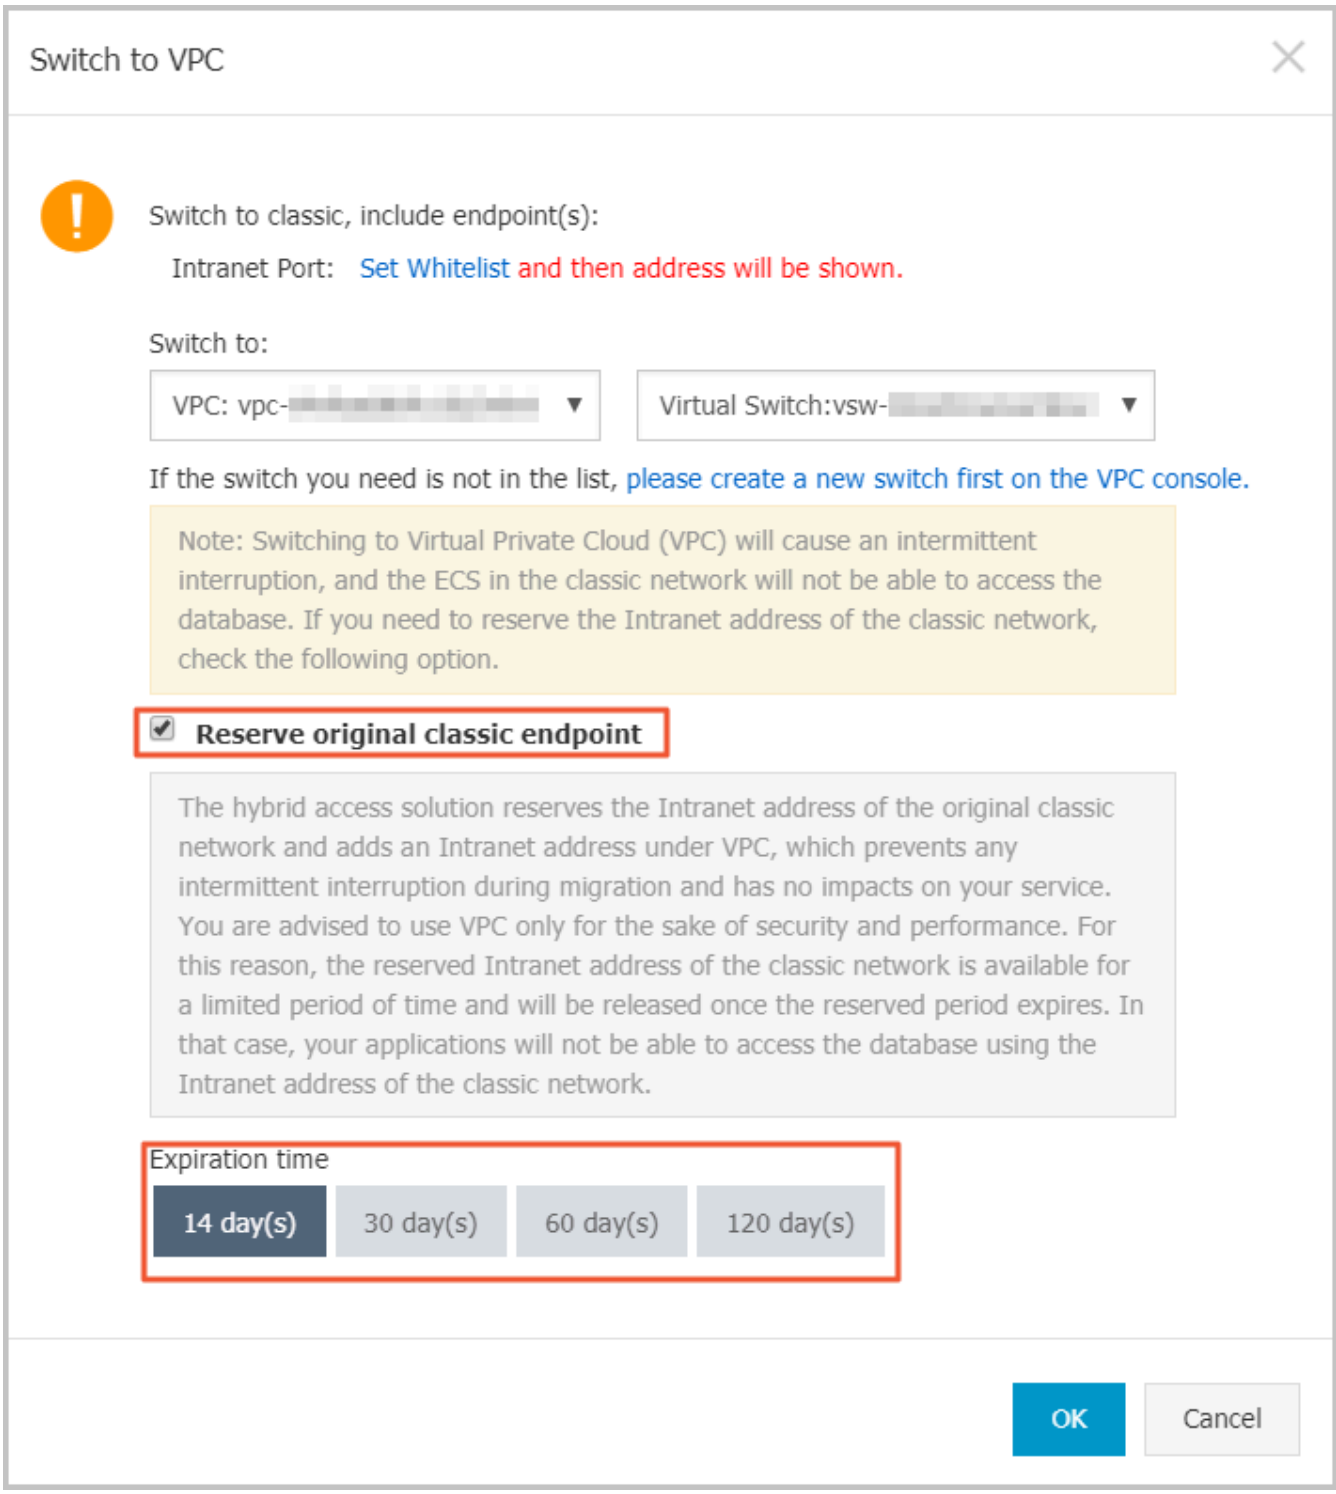

8. Click OK.

   The Original classic endpoint area is displayed, as shown in the following figure.

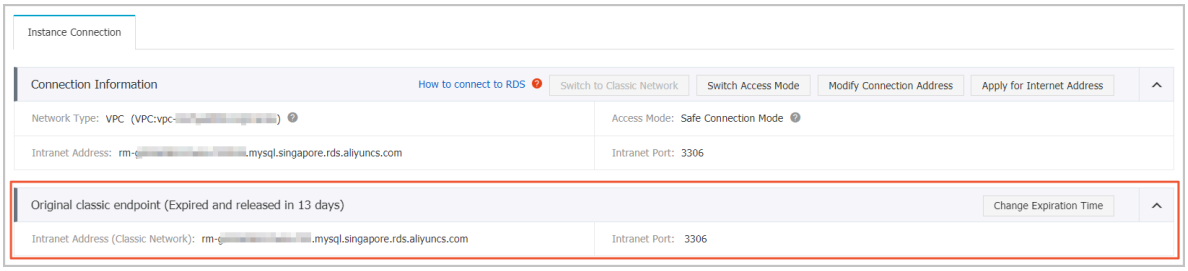

Change the expiration time of the original classic network

During the hybrid access period, you can change the reservation period of the intranet IP address of the original classic network at any time as needed, and the expiration date is recalculated from the new date. For example, if the intranet IP address of the original classic network is set to August 18, 2017, and you change the expiration time to 14 days later on August 15, 2017, the address is released on August 29, 2017.

Follow these steps to change the expiration time:

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. In the left-side navigation pane, click Connection Options to enter the Connection Options page.
5. On the Instance Connection tab page, click Change Expiration time, as shown in the following figure.

6. On the Change Expiration Time confirmation page, select an expiration time and click OK.

9.4 Set intranet and Internet IP addresses

You can select the connection type (intranet or Internet) of the instance according to your business requirements. The system generates an intranet IP address by default , so this document mainly introduces how to apply for an Internet IP address, set the connection address of the Internet or intranet, and release an Internet IP address.

Background information

RDS supports connections through the both intranet and Internet. The *series*, version, and *access mode* have the following effects on the selection of the connection address.

Instance series	Instance version	Access mode	Connection address
Basic Edition	· MySQL 5.7 · SQL Server 2012	Standard mode	· Intranet IP address · Internet IP address · intranet and Internet IP addresses
High-availability Edition	· MySQL 5.5/5.6 · SQL Server 2008 R2 · PostgreSQL 9.4 · PPAS 9.3	Standard mode	· Intranet IP address · Internet IP address
		Safe connection mode	· Intranet IP address · Internet IP address · intranet and Internet IP addresses
Finance Edition	MySQL 5.6	Standard mode	· Intranet IP address · Internet IP address
		Safe connection mode	· Intranet IP address · Internet IP address · intranet and Internet IP addresses

The applicable scenarios of the connection addresses are as follows:

- Use the intranet IP address only:
  - The system provides an intranet IP address by default and you can directly modify the connection address.
  - This scenario is applicable when your application is deployed on the ECS instance that is located in the same region and has the same network type as your RDS instance.
- Use the Internet IP address only:
  - This scenario is applicable when your application is deployed on the ECS instance that is located in the different region from that of your RDS instance.
  - This scenario is applicable when your application is deployed on a platform other than Alibaba Cloud.
- Use both of the intranet and Internet IP addresses:
  - This scenario is applicable when your application is deployed on the ECS instance that is located in the same region and has the same *network type* as your RDS instance, and application modules are deployed in an ECS where your RDS instance is not located.
  - This scenario is applicable when your application is deployed on the ECS instance that is located in the same region and has the same *network type* as your RDS instance, and on a platform other than Alibaba Cloud.

Attentions

- Before accessing the database, you must add the IP addresses or IP address segments that are allowed to access the database to a whitelist. For more information, see *Set the whitelist*.
- Traffic fees are charged for connections through Internet. For more information about pricing and fees charging, see *RDS Pricing*.
- Connecting the RDS instance through an Internet IP address may reduce the instance security. Proceed with caution. To get a higher transmission rate and a higher security level, we recommend that you migrate your applications to an ECS instance that is in the same region as your RDS.

Apply for an Internet IP address

1. Log on to the *RDS console* .
2. Select the region where the target instance is located.

3. Click the ID of the instance to visit the Basic Information page.
4. Click Connection options in the left-side navigation pane.
5. Click Apply for Internet Address, as shown in the following picture.

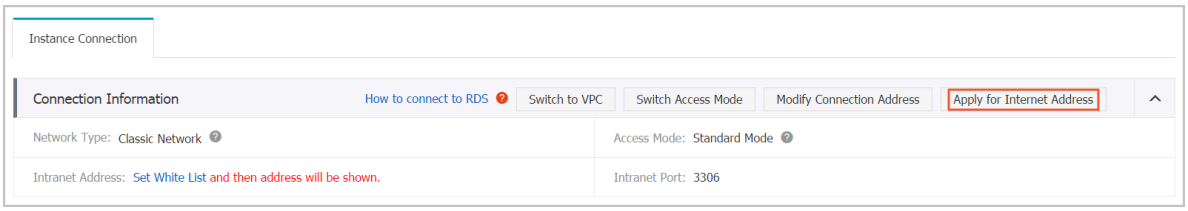

6. On the displayed confirmation window, click OK to generate an Internet IP address.

Modify the connection address

You can modify the Internet and intranet connection address based on your needs.

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. Click Connection options in the left-side navigation pane.
5. Click the Instance Connection tab.
6. In the Connection Informatio n area, click Modify Connection Address.

7. Select the connection type and modify its connection addresses and port number, and then click OK, as shown in the following figure.

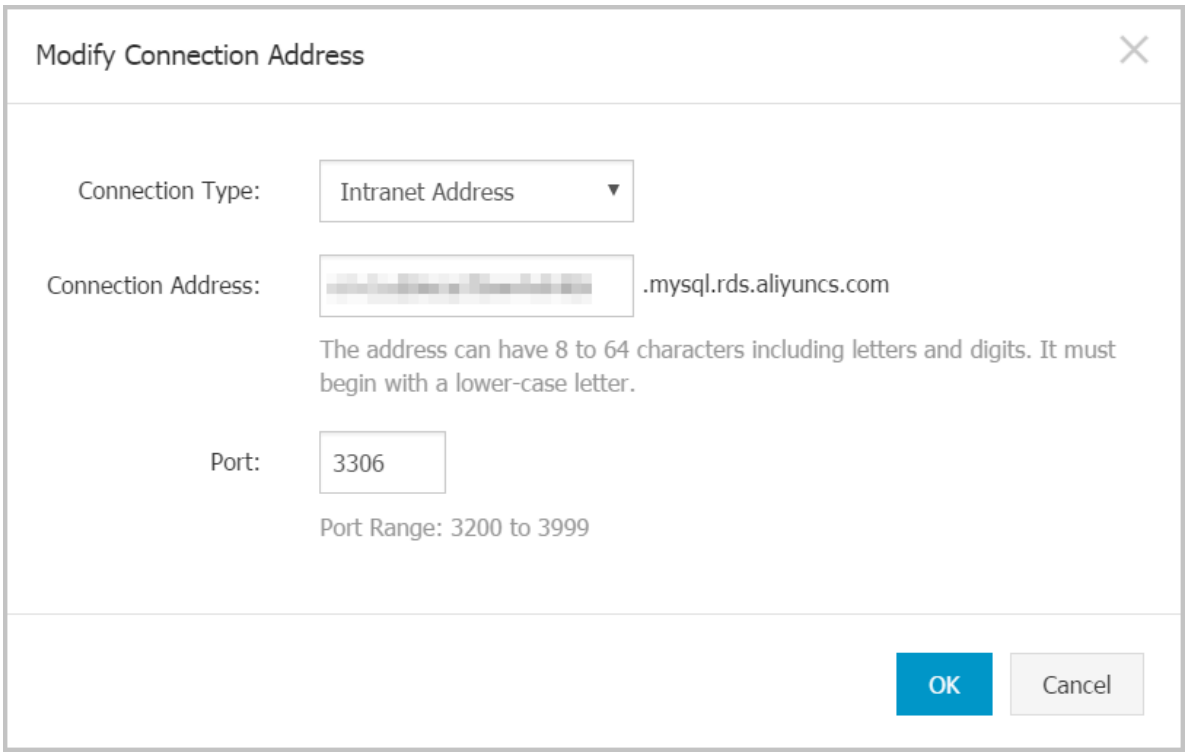

Parameters description:

- `Connection Type`: Select `intranet address` or `Internet address` according to the connection type to be modified.
- `Connection Address`: The address format is `xxx.sqlserver.rds.aliyuncs.com` and `xxx` is a user-defined field. The address contains 8 to 64 characters including letters and digits. It must begin with a lower-case letter.
- `Port`: indicates the number of the port through which RDS provides external services, which can be an integer within the range [3200, 3999].

Release an Internet IP address

If you want to release an Internet IP address, do as follows:

> **Note:**
> The operation can be performed only in `safe connection mode`. For more information about the safe connection mode, see *Disable database proxy Mode*.

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.

3. Click the ID of the instance to visit the Basic Information page.
4. Click Connection options in the left-side navigation pane.
5. Click the Instance Connection tab.
6. In the Connection Information area, click Release Internet Address.

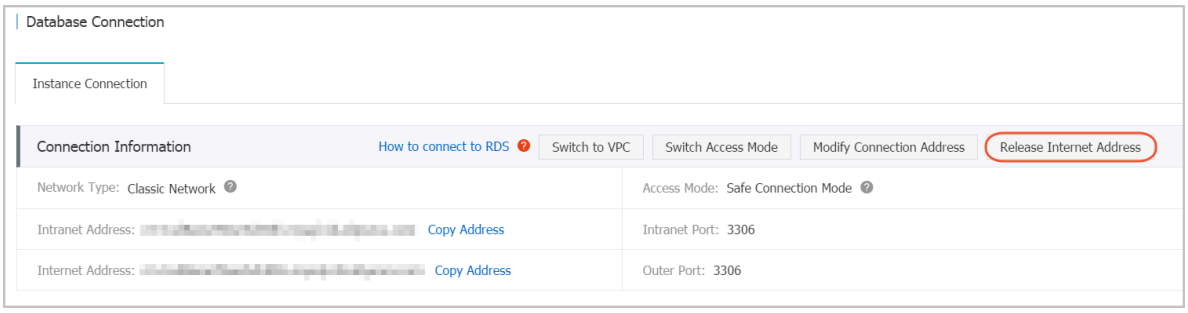

7. Click Confirm on the displayed confirmation dialog box to release the Internet IP address.

10 Monitoring and Alarming

10.1 Set the monitoring frequency

Background information

The RDS console provides abundant performance metrics for you to conveniently view and know the running status of instances. You can use the RDS console to set the monitoring frequency, view monitoring data of a specific instance, create monitoring views, and compare instances of the same type under the same account.

Two monitoring frequencies provided before May 15, 2018

- Once per 60 seconds (monitoring period: 30 days)
- Once per 300 seconds (monitoring period: 30 days)

Second-level monitoring frequency introduced since May 15, 2018

Minute-level monitoring frequencies cannot meet monitoring requirements of some users and maintenance personnel. Therefore, since May 15, 2018, RDS has introduced second-level monitoring frequencies. This facilitates problem locating and improves customer satisfaction.

- Once per 5 seconds (monitoring period: 7 days), turning to once per minute since the eighth day
- The detailed monitoring policies are described in the following table.

Instance type	Once per 5 seconds	Once per minute (60 seconds)	Once per 5 minutes (300 seconds)
Basic Edition	Not supported	Supported for free	Default configuration
High-availability or Finance Edition: Memory < 8 GB	Not supported	Supported for free	Default configuration
High-availability or Finance Edition: Memory >= 8 GB	Supported (Not free)	Default configuration	Supported for free

Restrictions

- You can configure second-level monitoring for instances that meet the following conditions:
  - The instance is located in these regions: China (Hangzhou), China (Shanghai), China (Qingdao), China (Beijing), or China (Shenzhen)
  - The instance is an RDS for MySQL instance.
  - The instance storage type is local SSD.
  - The instance memory space is 8 GB or more.
- All engines (MySQL, SQL Server, ProstgraSQL, and PPAS) and database versions support the following monitoring frequencies:
  - Once per 60 seconds
  - Once per 300 seconds

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to enter the Basic Information page.
4. Click Monitoring and Alarms in the left-side navigation pane.

   Note:
   Different types of databases support different metrics. For more information, see List of monitoring items at the end of this document.
5. Click the Monitoring tab.
6. Click Set Monitoring Frequency.
7. Select the monitoring frequency in the Set Monitoring Frequency dialog box and click OK.

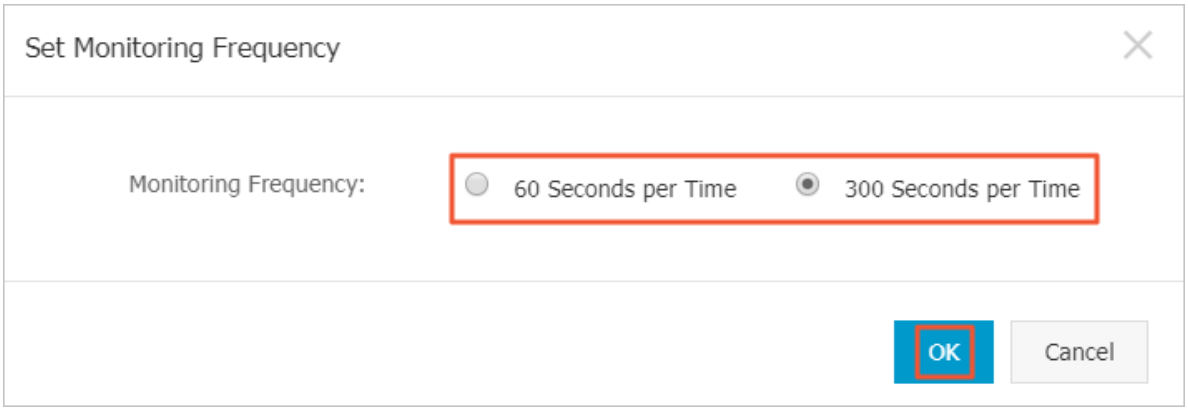

8. In the displayed Confirm dialog box, click OK.
9. On the Monitoring page, perform the following operations:

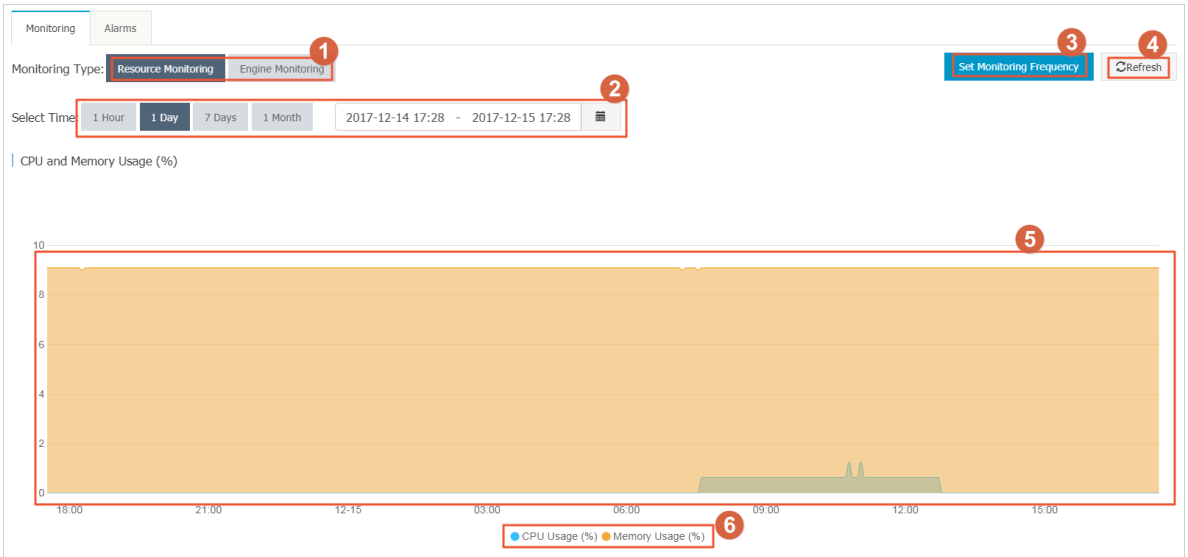

Interface description:

No.	Description
1	Select the monitoring type.
2	Select the monitoring period.
3	Set the monitoring frequency.
4	Refresh monitoring results.
5	View monitoring results.
6	Select monitoring items.

List of monitoring items

RDS for MySQL/MariaDB

Monitoring items	Description
Disk Space	Disk space usage of the instance, including: · Overall usage of the disk space · Data space usage · Log space usage · Temporary file space usage · System file space usage Unit: MB
IOPS	Number of I/O request times of an instance per second. Unit: time/second

Monitoring items	Description
Total Connections	Total number of current connections, including the number of active connections and total connections
CPU and Memory Usage	CPU usage and memory usage of an instance (excluding the memory used by OS)
Network Traffic	Incoming/outgoing traffic of an instance per second. Unit: KB
QPS/TPS	Number of SQL statements executed and transactions processed per second
InnoDB Buffer Pool	InnoDB buffer pool read hit rate, utilization rate, and percentage of dirty data blocks
InnoDB Read/Write Volume	Average InnoDB data read and write times per second. Unit: KB
Number of InnoDB Read and Write Times Per Second	Number of read and write times per second of InnoDB
InnoDB Log	Number of InnoDB physical writes to a log file, log write requests, and FSYNC writes to a log file per second
Temporary Tables	Number of temporary tables created automatically on the hard disk when the database executes SQL statements
MyISAM Key Buffer	Average key buffer read hit rate, write hit rate, and usage per seconcd of MyISAM
MyISAM Read and Write Times	Number of MyISAM read and write times from/to the buffer pool and from/to the hard disk per second
COMDML	Number of statements executed on the database per second. The statements include: · `Insert` · `Delete` · `Insert_Sel ect` · `Replace` · `Replace_Se lect` · `Select` · `Update`

Monitoring items	Description
ROWDML	Number of operations performed on InnoDB, including: · Number of physical writes to a log file per second · Number of rows read in InnoDB tables per second · Number of rows updated, deleted, and inserted in InnoDB tables per second

RDS for SQL Server

Monitoring items	Description
Disk Space	Disk space usage of the instance, including: · Overall usage of the disk space · Data space usage · Log space usage · Temporary file space usage · System file space usage Unit: MB
IOPS	Number of I/O request times of an instance per second. Unit: time/second
Connections	Total number of current connections, including the number of active connections and total connections
CPU usage	CPU usage (including CPU used by OS) of an instance
Network traffic	Incoming/outgoing traffic of an instance per second. Unit: KB
TPS	Number of transactions processed per second
QPS	Number of SQL statements executed per second
Cache hit rate	Read hit rate of the buffer pool
Average full table scans per second	Average number of full table scan times per second
SQL compilations per second	Number of compiled SQL statements per second
Page writes of the checking point per second	Number of page write times of the checking point in an instance per second
Logons per second	Number of logons per second
Lock timeouts per second	Number of lock expiration times per second
Deadlocks per second	Number of deadlocks in an instance per second

Monitoring items	Description
Lock waits per second	Number of lock waiting times per second

RDS for PostgreSQL

Monitoring item	Description
Disk Space	Usage of the instance disk space. Unit: MB
IOPS	Number of I/O request times of the data disk and log disk in an instance per second. Unit: time/second

RDS for PPAS

Monitoring item	Description
Disk Space	Usage of the instance disk space. Unit: MB
IOPS	Number of I/O request times of the data disk and log disk in an instance per second. Unit: time/second

10.2 Set monitoring rules

RDS offers the instance monitoring function, and sends messages to you after detecting an exception in an instance. In addition, when the instance is locked due to the insufficient disk space, the system sends a message to you.

Background information

Alibaba CloudMonitor offers monitoring and alarming. CloudMonitor helps you set alarm rules for metrics. You must add alarm contacts while set a contact group. The alarm contacts and the contact group are notified immediately when an alarm is triggered in the event of exceptions. You can create an alarm contact group using a related metric.

Procedure

1. Log on to the *RDS console* .
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. Click Monitoring and Alarms in the left-side navigation pane.
5. Click the Alarms tab.

6. Click Set Alarm Rules to open the CloudMonitor console.

   Note:
   You can click Refresh to manually refresh the current status of the alarm metric.

7. Select Alarms > > Alarm Contacts in the left-side navigation pane to open the Alarm Contact Management page.

   Note:
   When alarm rules are set for the first time, if the alarm notification object is not a contact of the Alibaba Cloud account of RDS, the alarm contact and alarm contact group must be created first. If you have already set the alarm contact and the alarm contact group, go to Step 10.

8. Click Create Alarm Contact.
9. Enter the alarm contact information in the Set Alarm Contact dialog box, click Send verification code, enter the verification code sent to your mailbox, and click Save.

   Note:

   - We recommend that you perform the next step to create the alarm contact group after you add all alarm notification objects.
   - Click Edit to modify a contact, or click Delete to delete a contact.

10. On the Alarm Contact Management page, click the Alarm Contact Group tab.
11. Click Create Alarm Contact Group.
12. Fill in `Group Name` and `Descriptio n` , select a contact from Existing Contacts, click to add the contact to `Selected Contacts` , and click OK.

    Note:
    On the Alarm Contact Group page, you can click to modify a contact group, click X to delete a contact group, or click Delete to delete a contact in the contact group.

13. After creating the alarm contact group, choose Cloud Service Monitoring > ApsaraDB for RDS from the left-side navigation pane.
14. Select the region of RDS for which the alarm rule is to be set.
15. Find the target instance and click Alarm Rules in the Actions column.

    The system displays the metrics of the current alarm.
16. Click Create Alarm Rule to add new alarm rules.

Note:

You can click Modify, Disable, or Delete for the metrics as needed.

11 Security

11.1 SQL audit

The SQL audit function allows you to view SQL details and periodically audit RDS instances.

Attentions

- Certain RDS instance types do not support the SQL audit function.
- The SQL audit function does not affect instance performance.
- SQL audit logs are kept for 30 days.
- Exported SQL audit files are kept for 2 days.
- The SQL audit function is disabled by default. Enabling this function incurs charges. For more information, see *Pricing*.

Differences between SQL audit logs and binlog

For MySQL instances, you can use SQL audit logs or binlog to view incremental data. Differences between them are as follows:

- SQL audit logs: Similar to MySQL audit logs, SQL audit logs collect information about all DML and DDL operations. The information is obtained through network protocol analysis. The SQL audit function does not parse actual parameter values , and a small number of records may be lost when the SQL query volume is large. Therefore, using SQL audit logs to collect incremental data may be inaccurate.
- Binlog: Binary logs accurately record all ADD, DELETE, and MODIFY operations and can accurately recover incremental data. Binary logs are stored in the instance temporarily. The system regularly transfers them to OSS and they are stored on OSS for 7 days. The system cannot save binlog files where data is being written, so certain binary logs are not uploaded when you click Upload Binlog on the RDS console.

  Therefore, binary logs accurately record incremental data, but you cannot obtain real-time binary logs.

Enable SQL audit

1. Log on to the *RDS console*.

2. Select the region where the target instance is located.
3. Click the ID of the target instance to go to the Basic Information page.
4. In the left-side navigation pane, click Security.
5. Click the SQL Audit tab and click Enable now.

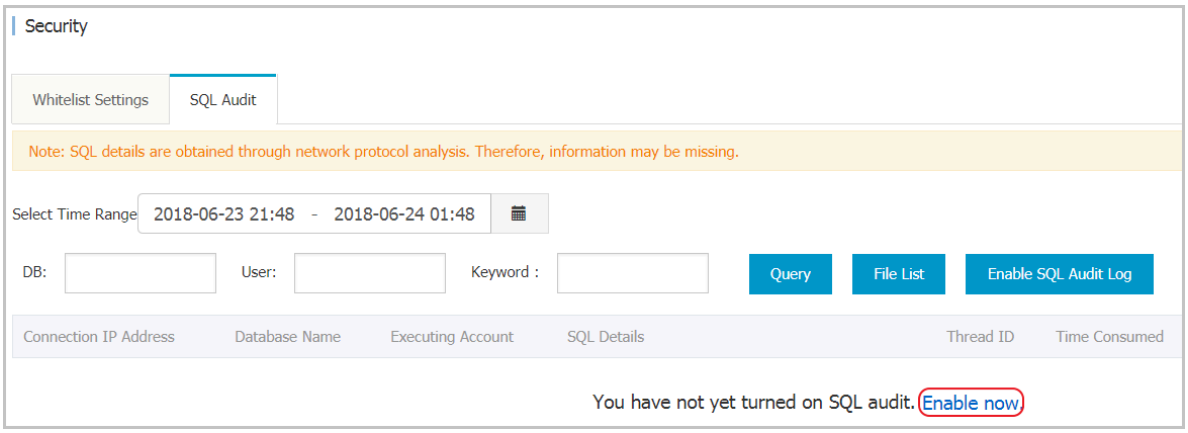

6. In the displayed dialog box, click Confirm.

Disable SQL audit

To save costs, you can disable the SQL audit function when you do not need it.

> **Note:**
> Disabling the SQL audit function deletes all SQL audit logs. Export logs before disabling the function.

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to go to the Basic Information page.
4. In the left-side navigation pane, click Security.
5. Click the SQL Audit tab. Click Export File and then click Confirm.

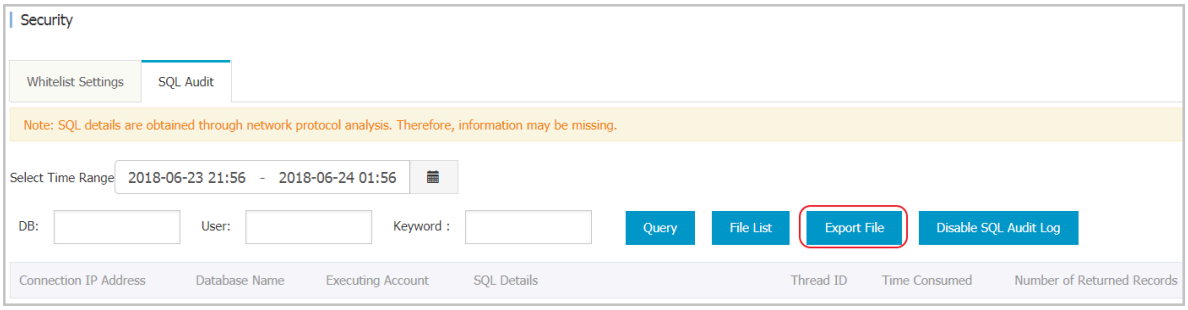

6. Download the SQL audit file and put it in a local directory.
7. Click Disable SQL Audit Log and then click Confirm.

11.2 Switch the IP whitelist to enhanced security mode

IP whitelist modes

RDS instances provide two IP whitelist modes:

- Standard mode: IP addresses in the whitelist apply to both classic networks and VPCs. This has security risks, so it is recommended that you switch to the enhanced security mode.
- Enhanced security mode: IP addresses in the whitelist are classified into two types: (1) IP addresses for classic networks and the Internet; (2) IP addresses for VPCs. In this mode, you need to specify the network type when you create an IP whitelist group.

  Currently, RDS for MySQL, PostgreSQL, and PPAS instances support the enhanced security mode.

Changes after switching to the enchanced security mode

- If the instance network type is VPC, a new whitelist group is generated and contains all IP addresses in the original whitelist. The new IP whitelist group applies only to VPCs.
- If the instance network type is classic network, a new whitelist group is generated and contains all IP addresses in the original whitelist. The new IP whitelist group applies only to classic networks.
- If the instance is in *hybrid access mode* (namely, an instance uses both a classic network and a VPC), two new whitelist groups are generated and each contain all IP addresses in the original whitelist. One of the whitelist group applies to VPCs and the other applies to classic networks.

> **Note:**
> The switch does not affect the *ECS security group* in the instance whitelist.

Attention

An IP whitelist can be switched from the standard mode to the enhanced security mode, and the switch is irreversible.

Procedure

1. Log on to the *RDS console*.

2. Select the region where the instance is located.
3. Click the ID of instance.
4. In the left-side navigation pane, select Security.
5. On the Whitelist Settings tab page, click Enable Enhanced Security Whitelist (Recommended).

6. In the displayed dialog box, click Confirm.

11.3 Set the whitelist

After an RDS instance is created, you need to set the whitelist so that servers can connect to the RDS instance. By default, the whitelist contains only the default IP address 127.0.0.1 and has no security group. This means that no server can access the RDS instance. The whitelist only controls access to the RDS instance and does not affect instance performance.

You can use either of the following methods to set the whitelist:

- Set the IP whitelist: Add IP addresses to the whitelist so that these IP addresses can access the RDS instance.
- Set the ECS security group: Add an ECS security group to the whitelist so that ECS instances in the security group can access the RDS instance.

We recommend that you periodically check and adjust the whitelist according to your requirements to maintain RDS security.

Attention

- The default IP whitelist group can only be modified or cleared, and cannot be deleted.

- % or 0.0.0.0/0 indicates that any IP address is allowed to access the RDS instance . This configuration greatly reduces the security of the database and is not recommended.
- If you cannot connect to the RDS instance after adding the application service IP address to the whitelist, you can obtain the actual IP address of the application by referring to *How to locate the local IP address using ApsaraDB for MySQL* .

Procedure

1. Log on to the *RDS console*.
2. In the upper left corner, select the region where the target instance is located.
3. Locate the target instance and click its ID.
4. In the left-side navigation pane, click Security to visit the Security page.
5. On the Whitelist Settings tab page, find the default whitelist group and click Modify.

   Note:
   You can also click Add a Whitelist Group to create a new group.

   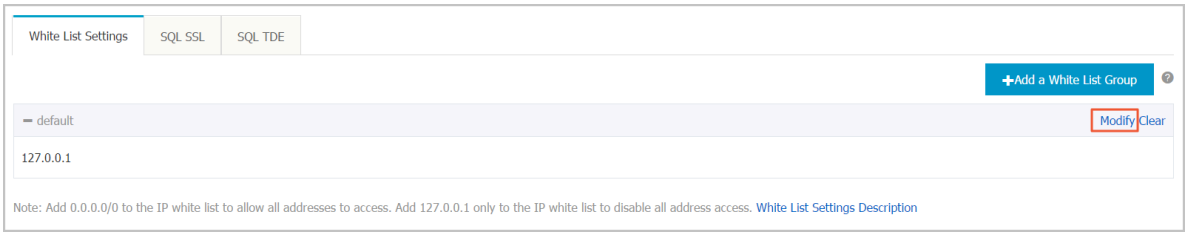

6. In the White List field of the displayed dialog box, add the IP addresses or IP address segments that need to access the RDS instance, and click OK.

   Note:

   - If you enter an IP address segment, such as 10.10.10.0/24, it indicates that any IP address in the format of 10.10.10.X can access the RDS instance.
   - If you want to enter multiple IP addresses or IP address segments, separate them by comma (but do not add blank spaces before or after commas), such as 192.168.0.1,172.16.213.9.

- If you click Upload ECS Intranet IP Address, the system displays the IP addresses of all ECS instances under your Alibaba Cloud account, and you can quickly add intranet IP addresses of ECS instances.

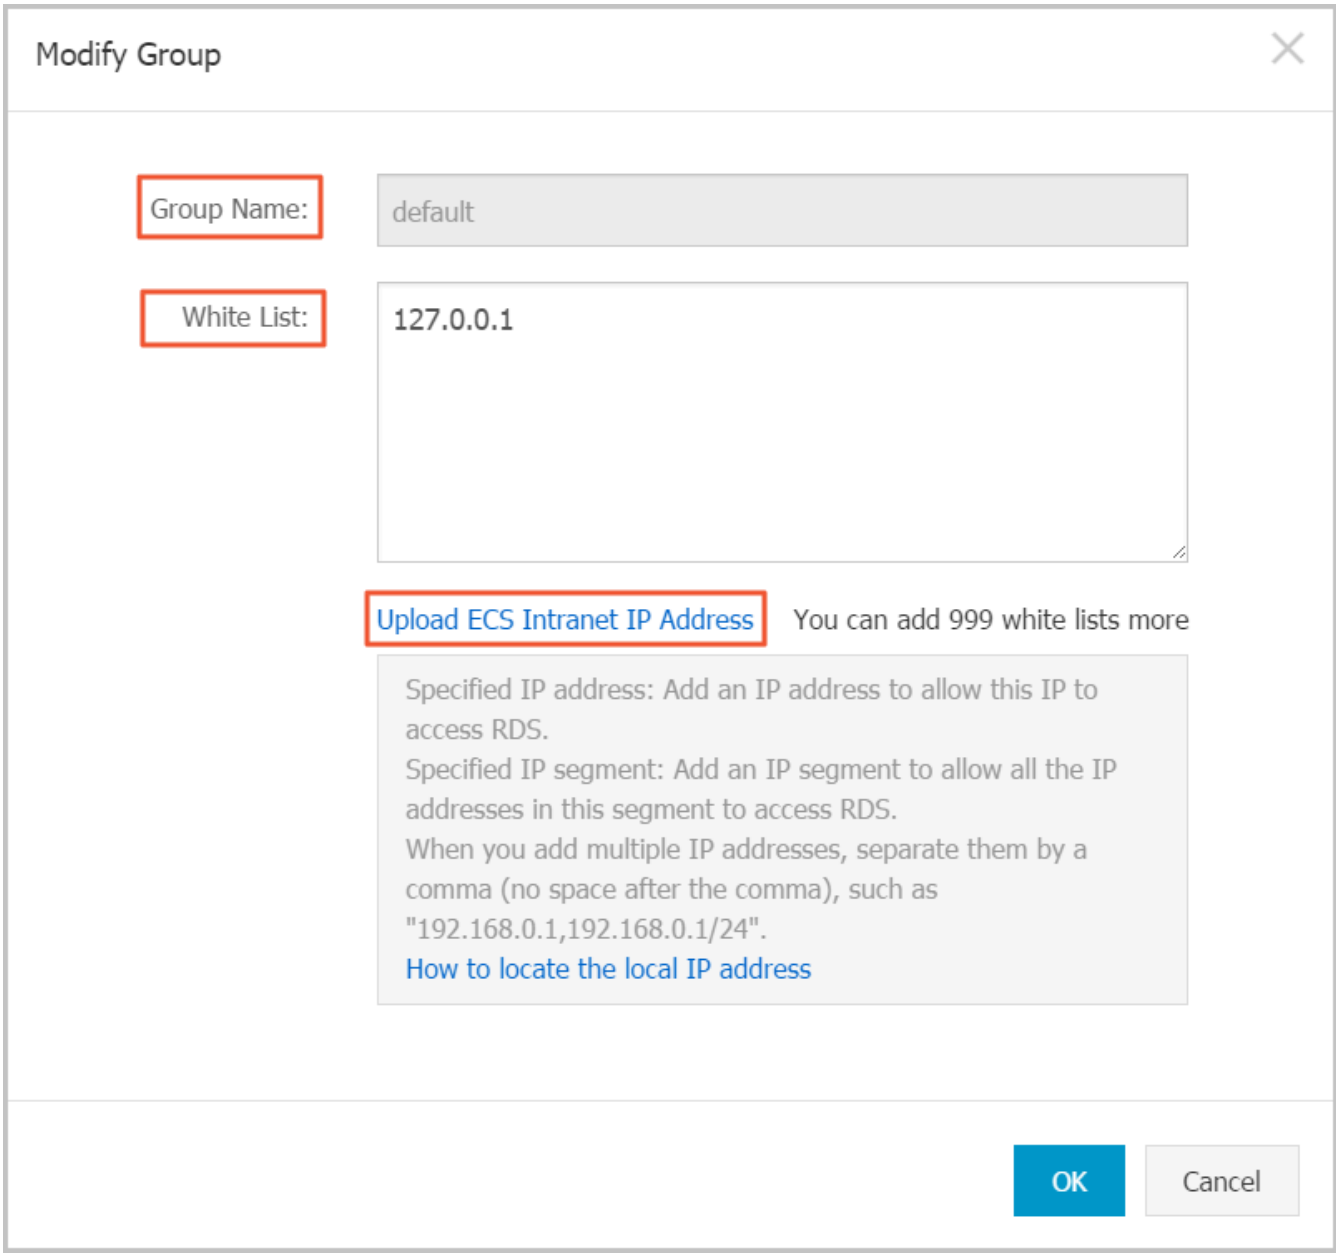

Add an ECS security group

A security group is a virtual firewall that is used to set network access control for one or more ECS instances. For more information about ECS security groups, see *Create a security group*.

Precautions

- RDS instances that support ECS security groups are MySQL 5.6, PostgreSQL, and MariaDB TX.

- Regions that support ECS security groups: Hangzhou, Qingdao, and Hongkong.
- You can set both the IP whitelist and ECS security group. All ECS instances specified in either the IP whitelist or security group can access the RDS instance.
- Currently each RDS instance supports one security group.

Procedure

1. Log on to the *RDS console*.
2. In the upper left corner, select the region where the target instance is located.
3. Locate the target instance and click its ID.
4. In the left-side navigation pane, click Security to visit the Security page.
5. On the Whitelist Settings tab page, click Add to Security Group.

   Note:
   Security groups marked with "VPC" are in VPCs.
6. Select a security group and click OK.

11.4 Set SSL encryption

To increase link security, you can enable Secure Sockets Layer (SSL) encryption and install an SSL certificate for necessary application services. SSL is used on the transport layer to encrypt network connections. It increases security and integrity of communication data, but also increases the network connection time.

Note:

- Due to the inherent drawbacks of SSL encryption, activating this function significantly increases your CPU usage. We recommend that you only enable SSL encryption for Internet connections requiring encryption. Intranet connections are relatively secure, and generally do not require link encryption.
- In addition, SSL encryption cannot be disabled once it is enabled. Therefore, enable SSL encryption with caution.

Enable SSL encryption

1. Log on to the *RDS Console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to enter the Basic Information page.

4. In the left-side navigation pane, click Security to go to the Security page.
5. Click the SSL tab.
6. Click the button next to Disabled, as shown in the following figure.

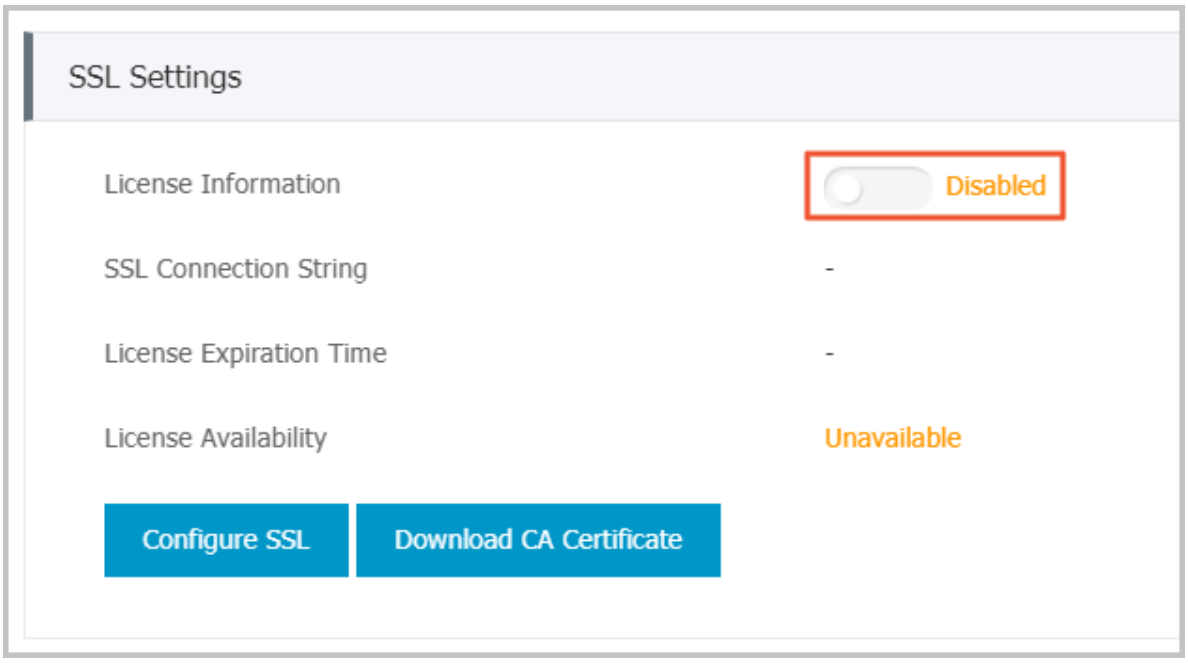

7. In the SSL Setting dialog box, select the link for which SSL encryption needs to be enabled and click OK to activate SSL encryption, as shown in the following figure.

Note:
You can choose to encrypt both Internet and intranet links as needed, but only one link can be encrypted.

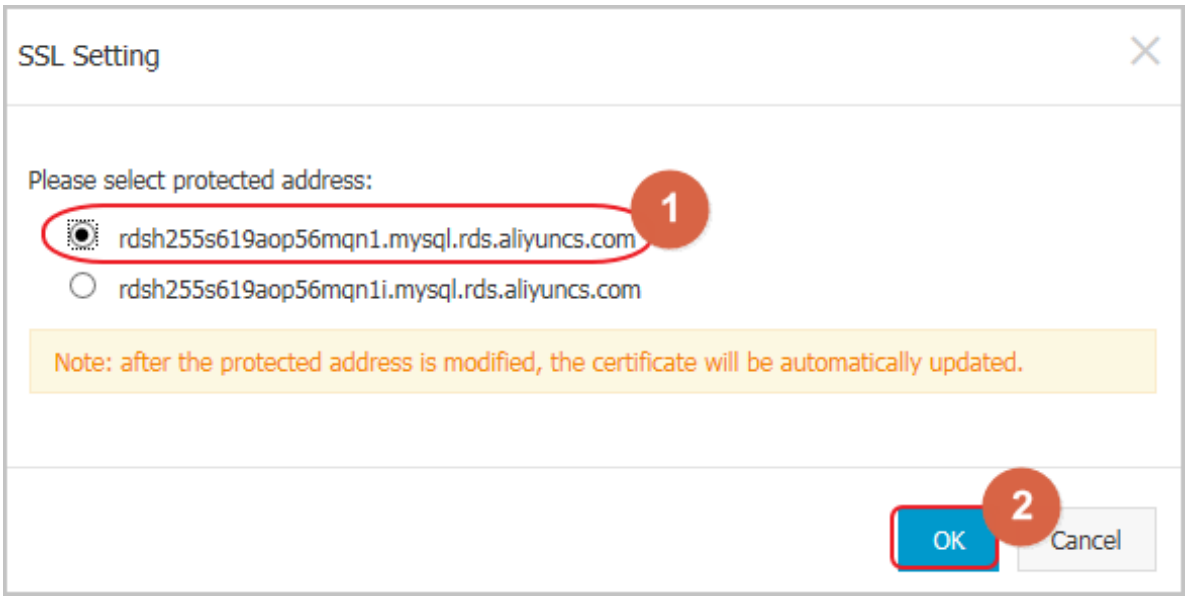

8. Click Download CA Certificate to download an SSL certificate, as shown in the following figure.

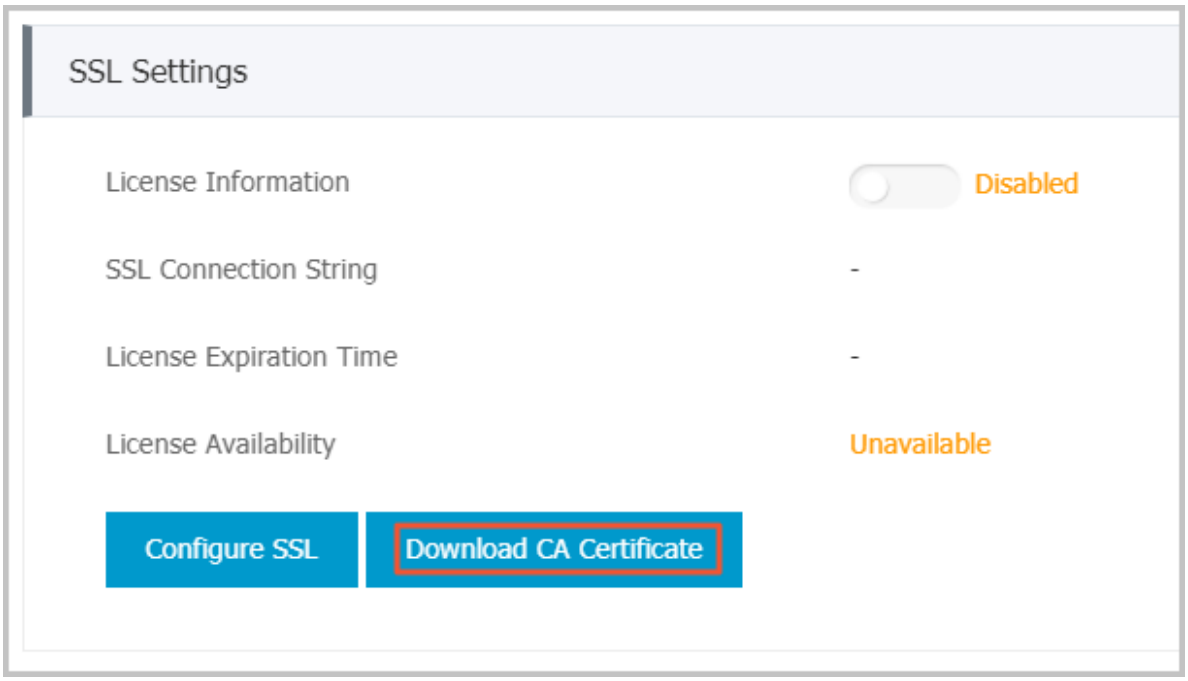

The downloaded SSL certificate is a package including the following files:

- p7b file: is used to import the CA certificate on Windows OS.
- PEM file: is used to import the CA certificate on other systems or for other applications.
- JKS file: is a Java truststore certificate file used for importing CA certificate chains in Java programs. The password is apsaradb.

> **Note:**
> When using JKS certificate files in Java, modify default jdk security configurations of jdk7 and jdk8 as follows: In the `jre / lib / security / java . security` file of the machine that runs the database to be accessed through SSL, modify the following configurations:

```
jdk . tls . disabledAl  gorithms = SSLv3 ,  RC4 ,  DH   keySize
<  224
```

```
jdk . certpath . disabledAl  gorithms = MD2 ,  RSA   keySize  <
1024
```

If you do not modify the JDK security configuration, the following error will be reported. Other similar errors are generally caused by Java security configurations.

```
javax . net . ssl . SSLHandsha  keExceptio  n :  DHPublicKe  y
does   not   comply   to   algorithm
           constraint  s
```

Configure the SSL CA certificate

After SSL encryption is enabled, you need to configure the SSL CA certificate for applications or clients that access RDS. The following uses MySQL Workbench as an example to describe how to install the SSL CA certificate. For other applications or clients, see their usage instructions.

1. Open MySQL Workbench.
2. Choose Database > Manage Connections .
3. Enable Use SSL and import the SSL CA certificate, as shown in the following figure.

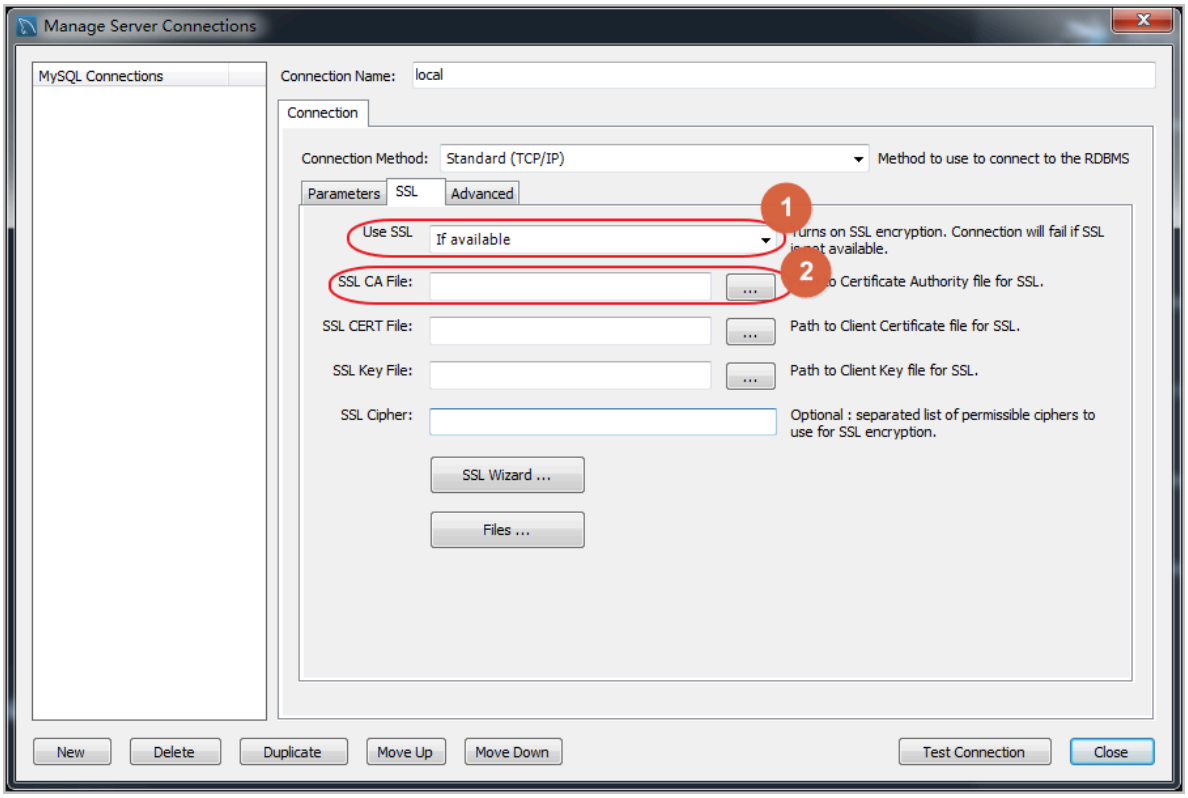

11.5 Set TDE

Transparent Data Encryption (TDE) can be used to perform real-time I/O encryption and decryption on instance data files. To improve data security, you can enable TDE to encrypt instance data.

Note:

Currently, TDE is only applicable to databases of SQL Server 2008 R2 and MySQL 5.6. To view or modify TDE settings, you need to log on with an Alibaba Cloud account rather than a RAM account.

Background information

TDE provides real-time I/O encryption and decryption on data files. The data is encrypted before being written to the disk and decrypted when being reading from the disk into the memory. TDE does not increase the size of data files. Developers do not have to modify any applications before using the TDE function.

Considerations

- Once TDE is activated, it cannot be deactivated.
- Encryption uses keys produced and managed by the Key Management Service (KMS). RDS does not provide the keys and certificates required for encryption. After TDE is activated, if you want to restore data to your local device, use RDS to decrypt the data first.
- After TDE is activated, CPU usage significantly increases.

Prerequisite

KMS is activated.

Procedure

1. Log on to the *RDS console* and select the target instance.
2. Click Data Security in the left-side navigation pane.
3. On the Data Security page, click the TDE tab.

4. Click Not Activated, as shown in the following figure.

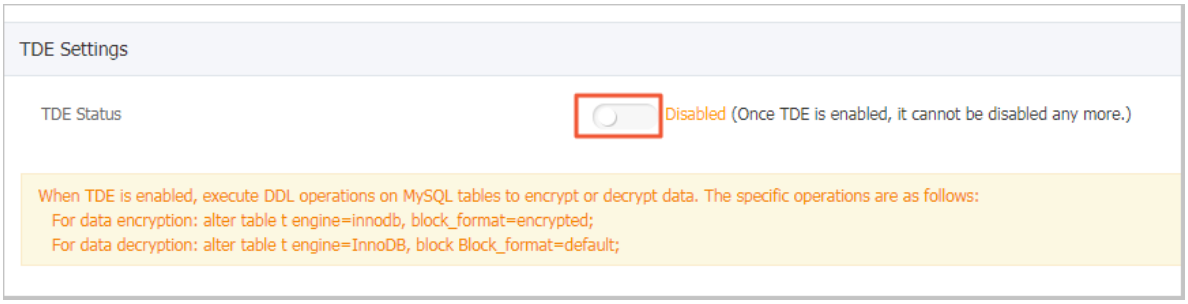

5. Click OK to activate TDE.

> **Note:**
> If you have not activated KMS, you are prompted to do so when activating TDE. After activating KMS, click Not Activated to activate TDE.

6. Log on to the database and run the following command to encrypt the relevant tables.

```
alter table <tablename> engine = innodb, block_form at =
encrypted ;
```

Subsequent operation

If you want to decrypt a table encrypted by TDE, run the following command.

```
alter table <tablename> engine = innodb, block_form at =
default ;
```

12 Log management

All instance versions except MySQL 5.7 support log management. You can use the RDS console or SQL statements to query error logs and slow SQL log details for fault analysis. However, you can manage logs of instances in SQL Server 2012 and later versions only through SQL statements. This document describes how to manage logs through the RDS console and SQL statements.

Use the RDS console to manage logs

You can use the RDS console to manage logs of MySQL 5.5/5.6, SQL Server 2008 R2, PostgreSQL, and PPAS instances. The actual interface may vary with engine types and versions.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to enter the Basic Information page.
4. Click Log Management in the left-side navigation pane.
5. On the Log Management page, select Error Log, Slow SQL Log Details, Slow SQL Log Summary, or Switch Logs, select a time range, and click Query.

Query item	Content
Error Log	Records the SQL statements that are failed to be executed in the past month.
Slow SQL Log Details	· Records the SQL statements that lasted for over one second (For MySQL and MariaDB, you can modify this time threshold by modifying the `long_query _time` parameter in Parameters) in the past month. Similar SQL statements are displayed once only. · The list does not include slow SQL logs of the past two hours. To query these logs, check the slow_log_view table in the MySQL database.
Slow SQL Log Summary	Provides statistics and analysis reports for SQL statements that lasted for over one second (For MySQL and MariaDB, you can modify this time threshold by modifying the `long_query _time` parameter in Parameters) in the past month.

Use SQL statements to manage logs

Instances in SQL Server 2012 and later versions read error logs only through the `sp_rds_rea d_error_lo gs` storage procedure. The method of using it is similar to that of using `sp_readerr orlog` .

Example 1:

```
EXEC sp_rds_rea d_error_lo gs
```

Example 2:

```
EXEC sp_rds_rea d_error_lo gs 0 , 1 ,' error '
```

13 SQL explorer

ApsaraDB RDS for MySQL has upgraded the SQL audit function as SQL explorer, which continues to provide security audit and performance diagnosis, but has more diverse features and costs much less. The upgrade process does not affect the RDS for MySQL instances.

Applicable scope

- RDS for MySQL 5.5
- RDS for MySQL 5.6
- RDS for MySQL 5.7 High-Availability Edition based on local SSDs

Upgrade plan

To ensure the quality of our services, all RDS for MySQL instances across the globe will be upgraded in several batches.

After the upgrade time, new and existing instances will both support the SQL explorer function.

Regions	Upgrade time
China (Hangzhou, Shanghai, Qingdao, Beijing, Shenzhen, Hong Kong)	By of end of December, 2018
Singapore, Malaysia, and Indonesia	By the end of January, 2019
Japan, Australia, India, and China (Zhangjiakou)	By the end of February, 2019
London	By the end of March, 2019
China (Hohhot), US (Silicon Valley, Virginia), and UAE (Dubai)	To be determined

Features

- SQL log: SQL log records all operations that have been performed on databases. With SQL log, you can do database troubleshooting, action analysis, and security audit.

- Enhanced search: You can search data by database, user, client ID, thread ID, execution time, or the number of scanned rows. You can also export and download the search results.

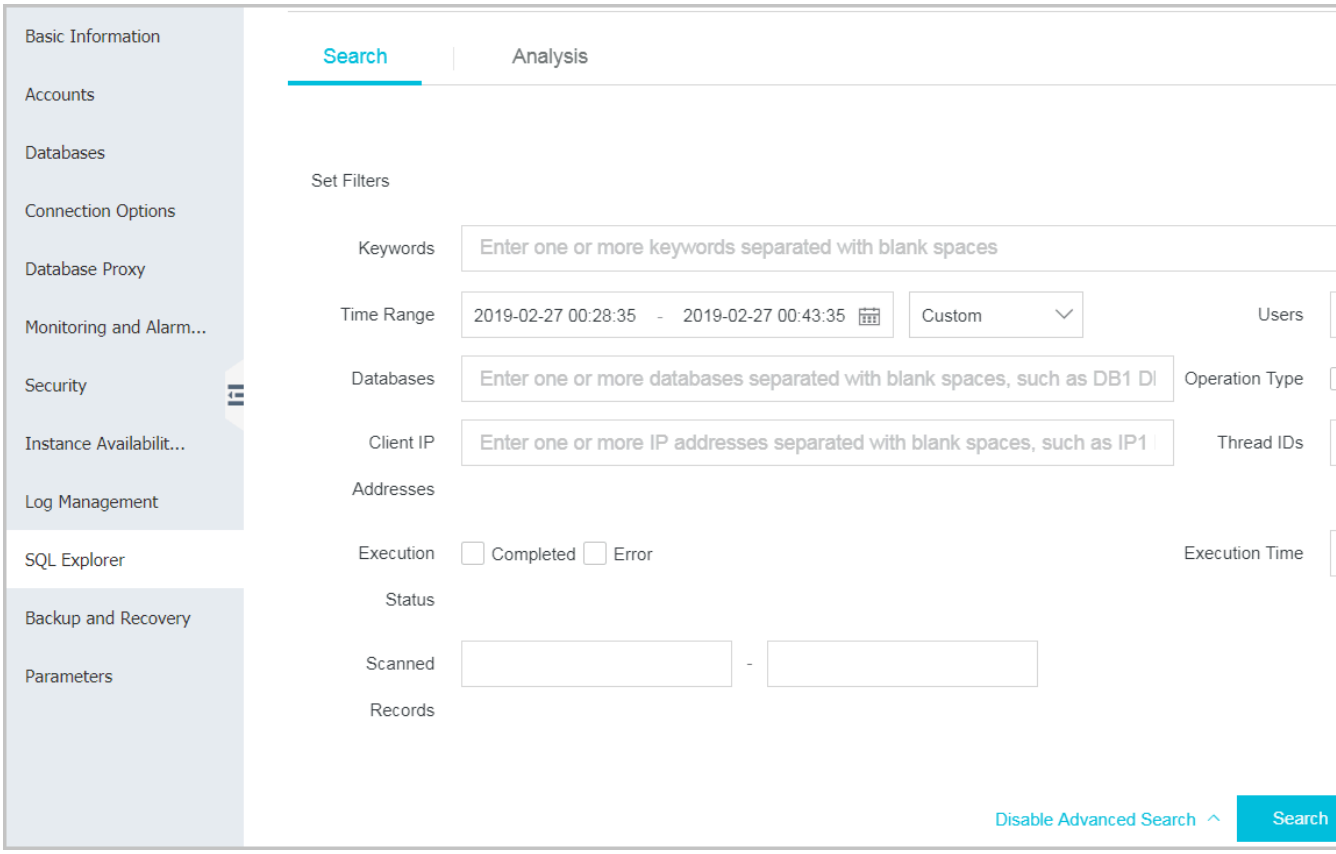

- SQL analysis: This new feature provides visualized interactive analysis of SQL log of a specified time period. You can use this feature to locate abnormal SQL statements and performance issues.
- Cost reduction: SQL explorer adopts the column-based storage and compression technology to reduce the SQL log size and reduce storage costs by about 60%. The hourly fee is US$ 0.0018 per GB.

Activate SQL explorer

1. Log on to the *RDS console*.
2. In the upper-left corner, select the region of the target instance.
3. Locate the target instance, and click the instance ID.
4. In the left-side navigation pane, select SQL Explorer.

5. Click Activate Now.

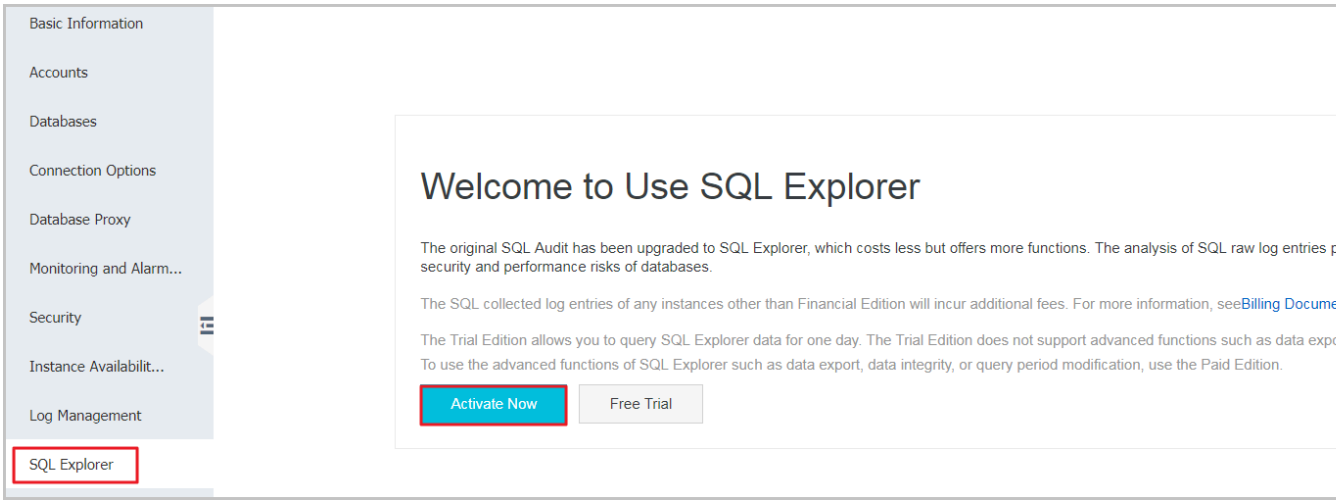

6. Specify the SQL log storage duration (for how long you want to keep the SQL log), and click Activate. The system then automatically starts charging an hourly fee of US$ 0.0018 per GB.

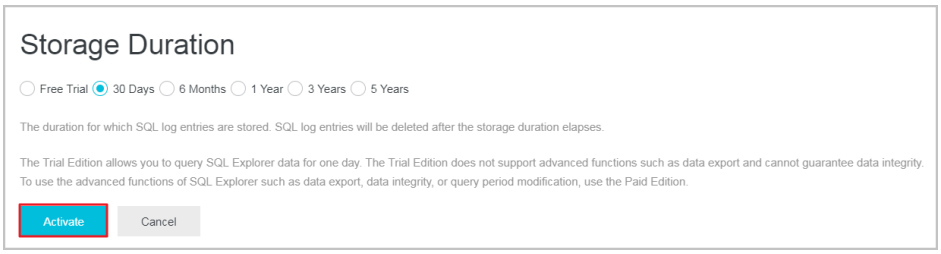

Modify the SQL log storage duration

1. Log on to the *RDS console*.
2. In the upper-left corner, select the region of the target instance.
3. Locate the target instance, and then click the instance ID.
4. In the left-side navigation pane, select SQL Explorer.
5. Click Service Settings.

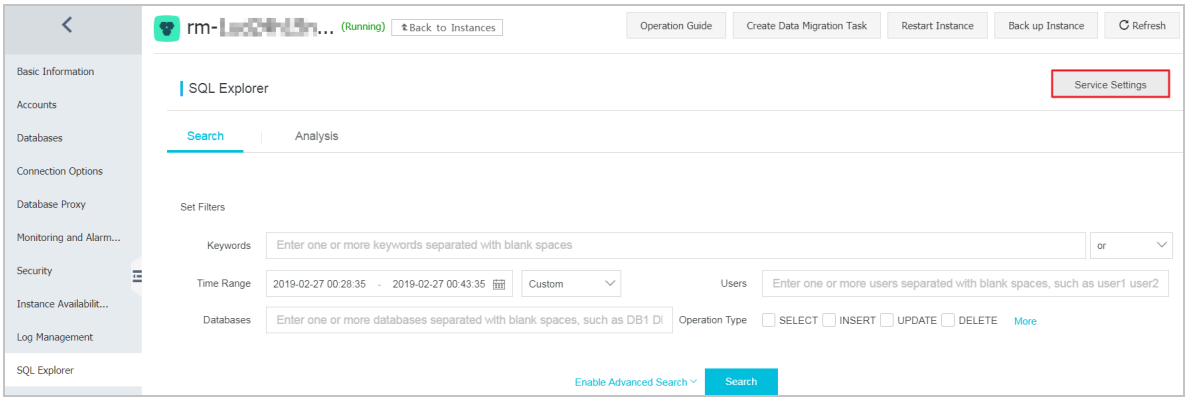

6. Modify the storage duration.

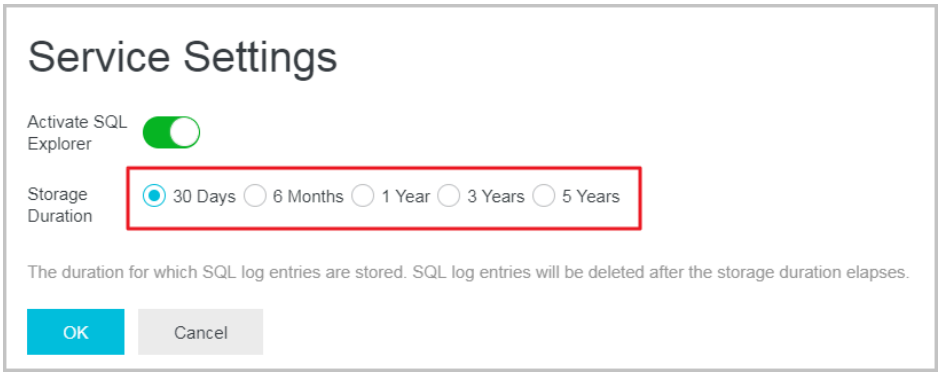

Disable SQL explorer

Note:

If you disable the SQL explorer function, the existing SQL log will be deleted. Export and save the SQL log to your local disks before you disable the function.

1. Log on to the *RDS console*.
2. In the upper-left corner, select the region of the target instance.
3. Locate the target instance, and then click the instance ID.
4. In the left-side navigation pane, select SQL Explorer.
5. Click Export.

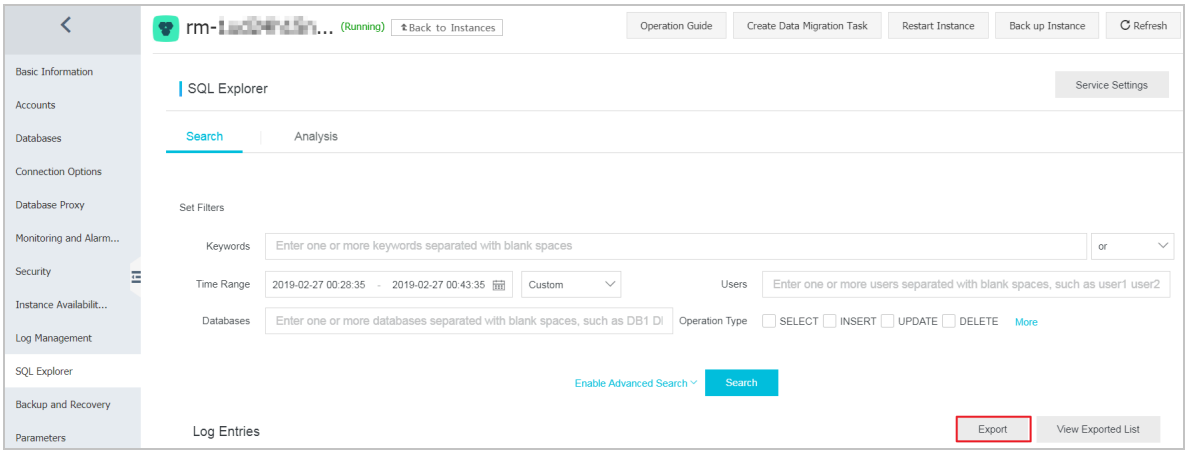

6. Click OK in the dialog box.

7. After the export process is complete, click View Exported List and then download the log file.

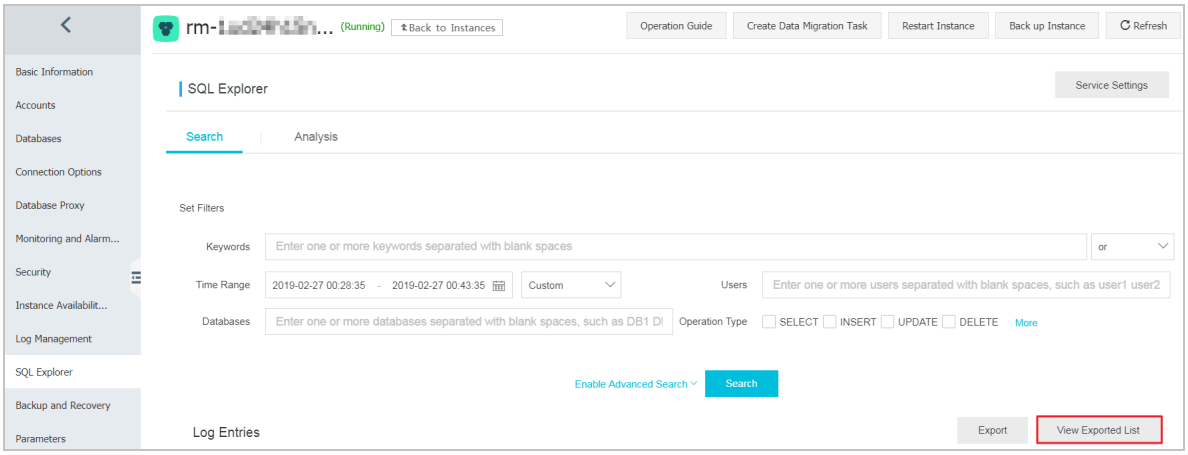

8. Click Service Settings.

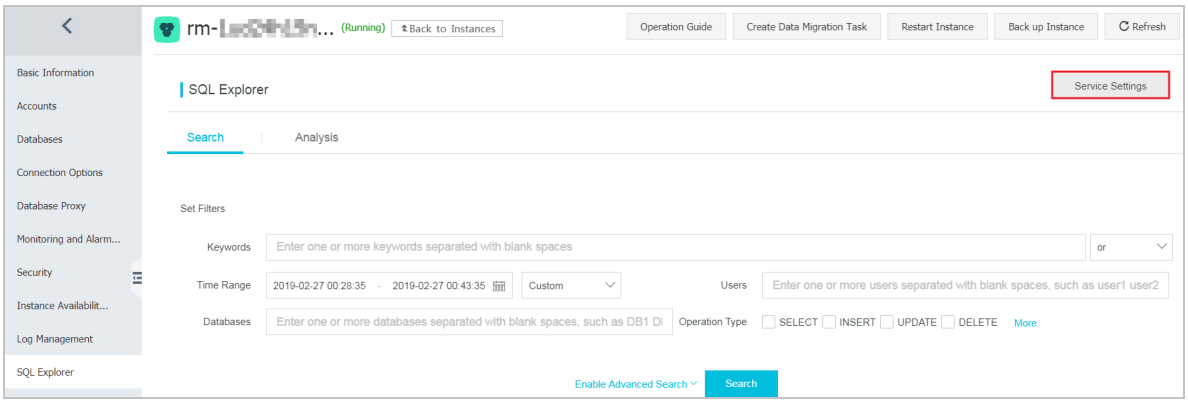

9. Click the toggle to disable the SQL explorer function.

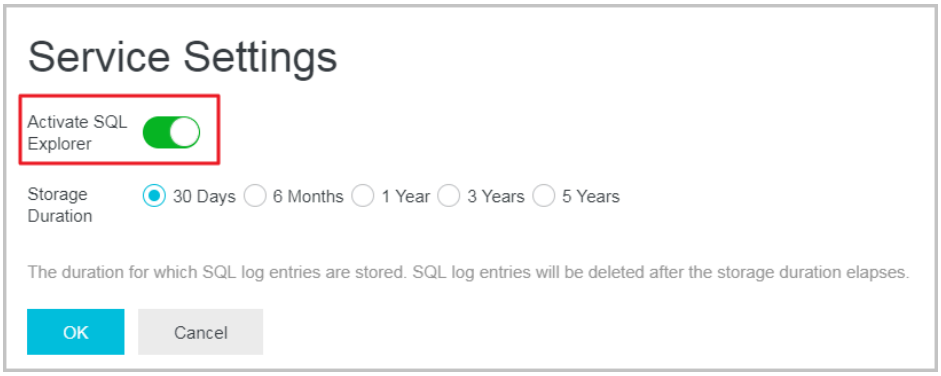

14 Backup

14.1 Back up RDS data

You can configure a backup policy to adjust the cycles of RDS data backup and log backup. As a result, RDS enables the auto-backup feature. You can also manually back up RDS data.

Instance backup files occupy backup space. Charges are incurred if the used space exceeds the free quota. You must set a backup cycle appropriately to cater to the service requirements based on the available backup space. For information about the free quota, see *View the free quota of the backup space*. To view the charging standard for backup space usage, see *Pricing*.

Backup policies

ApsaraDB supports data backup and log backup. To recover data to a point in time, you must enable the log backup function. The following table lists the backup policies applicable to different database types:

Database type	Data backup	Log backup
MySQL	· MySQL 5.5/5.6/5.7 (including High-availability Edition and Finance Edition): - Automatic backup supports full physical backup. - Manual backup supports full physical backup, full logical backup, and single-database logical backup. · MySQL 5.7 Basic Edition: - Supports only snapshot-based backup instead of logical backup. - Backup files are retained for at most 7 days for free.	· After being generated, binlogs (500 MB per log) are compressed and uploaded immediately. Local files are deleted within 24 hours. · Binlog files occupy instance disk capacity. Using the binlog upload function, you can upload binlog files to OSS. This does not affect the data recovery function and stops the binlog files from occupying instance disk space.

Database type	Data backup	Log backup
SQL Server	· Supports full physical backup and incremental physical backup. · Automatic backup cycles from full backup, incremental backup to incremental backup. For example, if a full backup is performed on Monday, incremental backups are performed on Tuesday and Wednesday, and another full backup is performed on Thursday, with incremental backups on Friday and Saturday. If a full backup is manually performed at any time in the backup cycle, the next two backups are incremental backups. · SQL Server always compresses transaction logs during the backup process. On the Backup and Recovery page of the target instance's management console, you can click Compress Transaction Log to manually compress transaction logs.	· RDS automatically generates log backups (log files). You can set the log file generation interval to 30 minutes or the data backup interval. The interval does not change the total size of generated log files. · The log backup function cannot be disabled. · You can set the log backup reservation duration to a time period ranging from 7 to 730 days. · You can download log files.
PostgreSQL	Supports full physical backup.	After being generated, write-ahead logs (WALs) (16 MB per log) are compressed and uploaded immediately. Local files are deleted within 24 hours.
PPAS	Supports full physical backup.	After being generated, WALs (16 MB per log) are compressed and uploaded immediately. Local files are deleted within 24 hours.

Configure automatic backup (Set backup policies)

> **Note:**
> The following uses MySQL 5.7 (High-availability Edition) as an example.

1. Log on to the *RDS console*.
2. Click the ID of the instance to visit the Basic Information page.
3. Click Backup and Recovery in the left-side navigation pane.
4. On the Backup and Recovery page, select Backup Settings and click Edit.

5. In the Backup Cycle dialog box, set backup parameters and click OK.

   The parameters are explained as follows:

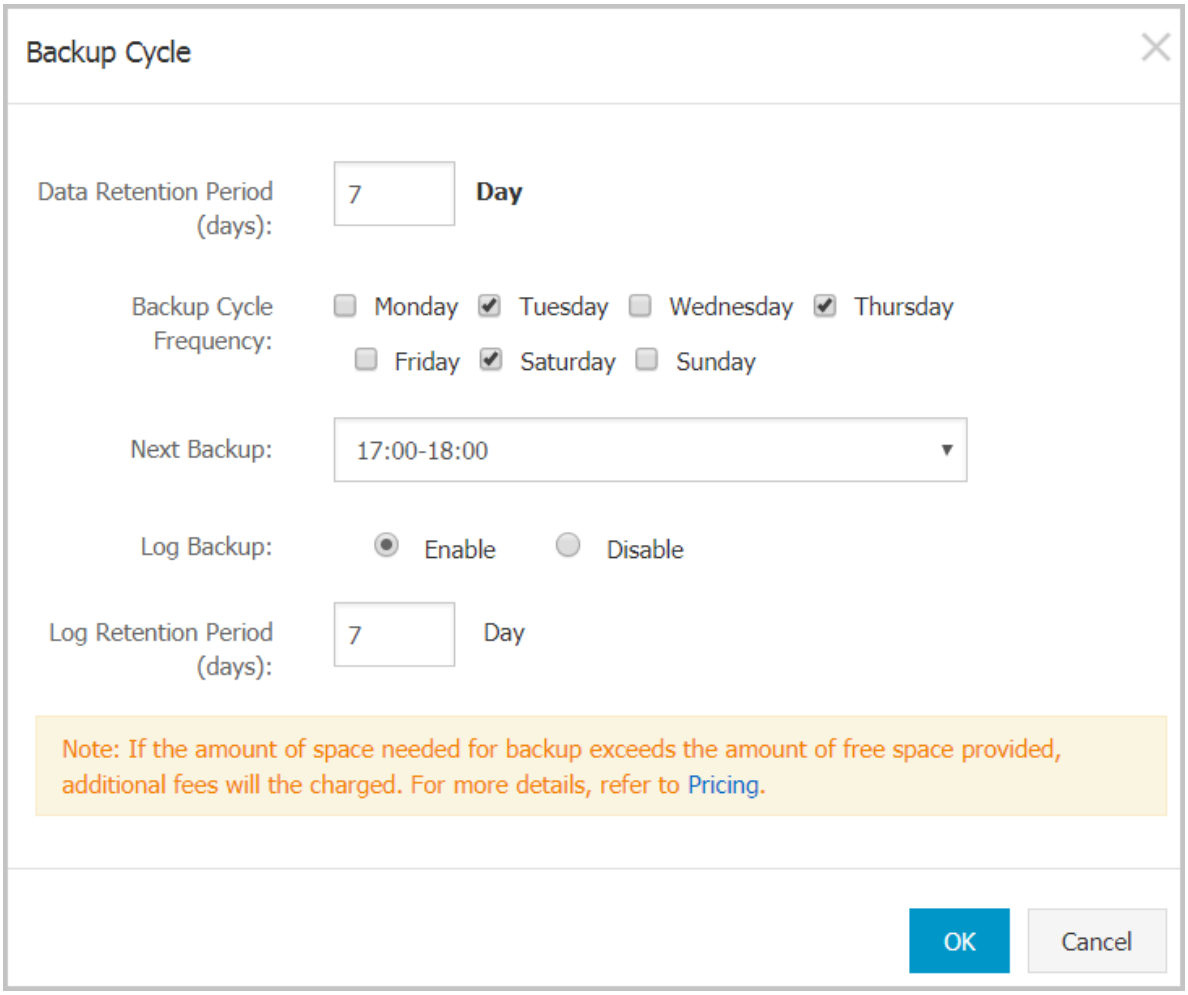

Parameters	Description
Data Retention Period (days)	· Specifies the time period during which backup files are retained. The default value is 7 days. The value range is 7 to 730 days. · MySQL 5.7 Basic Edition backup files are retained for free for at most 7 days.
Backup Cycle Frequency	· You can set it to one or multiple days in a week. · SQL Server, PostgreSQL, and PPAS instances are backed up every day by default, which cannot be modified.
Next Backup	This parameter can be set to any time. Units: Hour
Log Backup	Possible values are Enable and Disable.

Parameters	Description
Log Retention Period (days)	· Specifies the number of days during which log backup files are retained. The default value is 7 days. · The value range is 7 to 730 days and it must be less than or equal to the value of the retention days.

Manual backup

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. Click Back up Instance at the upper right corner.
5. Set Backup Mode and Backup Policy.

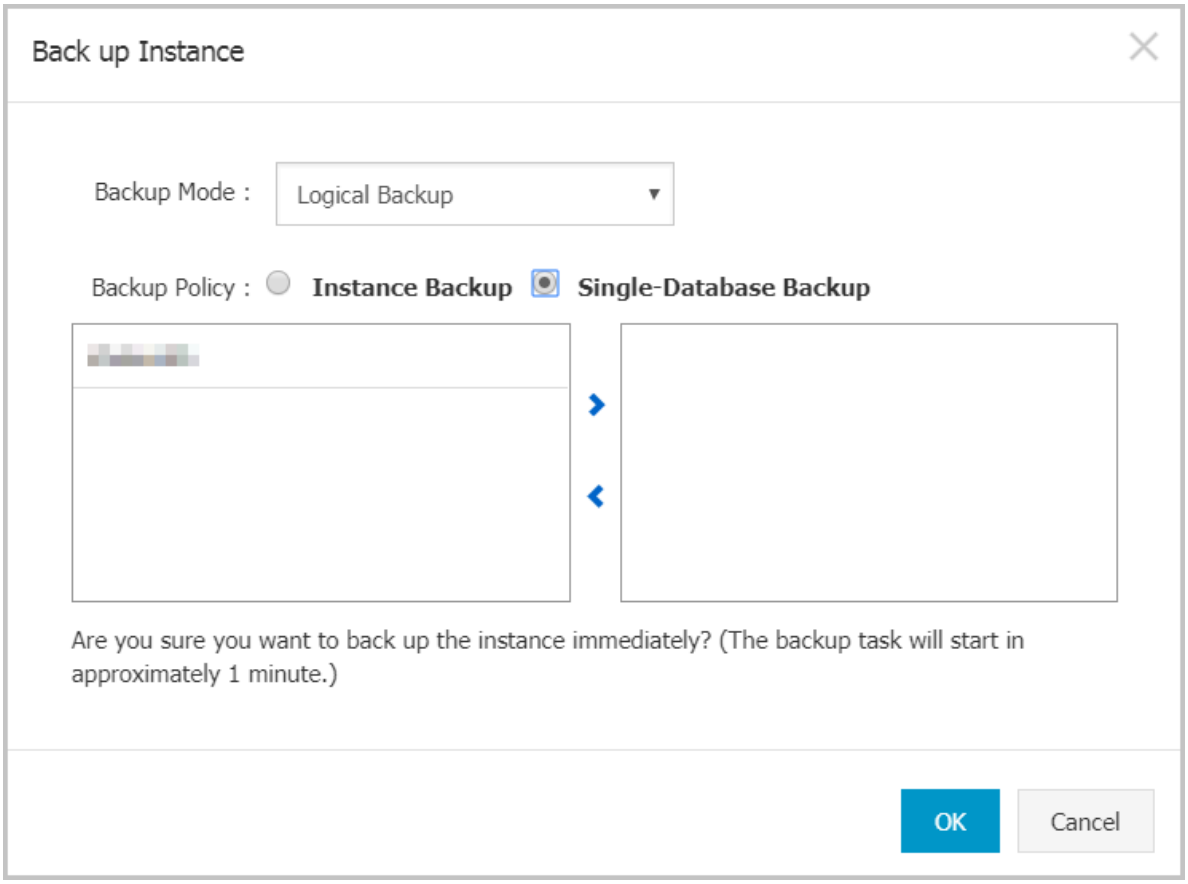

Note:

· The backup mode and policy vary with the database type. For more information, see *Backup policies*

- If you choose single-database backup, click > to select a database to be backed up. If you do not have a database, create one by referring to *Create a database*.

6. Click OK.

14.2 View the free quota of the backup space

Backup files of an instance occupy the backup space. Each RDS instance provides the backup space with a certain free quota. Additional charges can be incurred for the backup space exceeding the free quota. For information about billing standards for backup space usage, see *RDS pricing*. Different types of instances have different free backup space quotas. This document describes how to view and calculate the free quota of the instance backup space.

Formula for calculating the free quota of the backup space

If the total volume of your backup data (OSS and Archive Storage) and backup log (OSS) is less than or equal to 50% of the storage space bought for the instance, the space is within the free quota.

The excess backup space beyond the free quota is billed by hour. (Unit: GB, rounded up only)

```
Costs per hour = data backup volume + Log backup
volume - Instance storage space x 50 %
```

View the free quota of the backup space on the RDS console

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the target instance to go to the Basic Information page.
4. In the Resource Information area at the bottom of the page, check the remarks next to Backup Size, which shows the free quota, as in the following figure.

> **Note:**
> Instances of different types support different free quotas. The following figure is only an example.

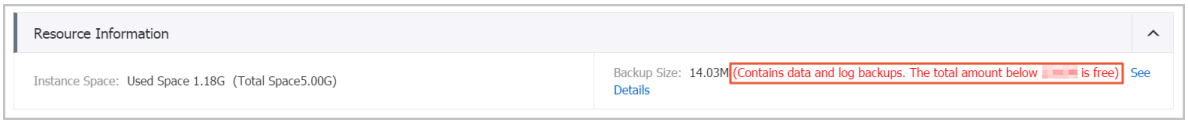

14.3 Download data and log backup files

You can download data and log backup files that are not encrypted.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. Click Backup and Recovery in the left-side navigation pane.
5. Do as follows to download a data or log backup file:
   - To download data backups, click the Backup List tab.
   - To download log backups:
     a. Click the Binlog List tab for MySQL and SQL Server.
     b. Click the Archive List tab for PostgreSQL and PPAS.
   - Specify a time range.
   - Find the data backup or log backup you want and click Download in the Action column.

Note:
If the binlog file is used for restoring data to an on-premise database, pay attention to the following:

- Instance Number of the binlog must be the same as Instance Number of the data backup.

- The binlog backup start time must be later than the data backup time and earlier than the restoration time.

· In the Download Instance Backup File dialog box, select a download method.

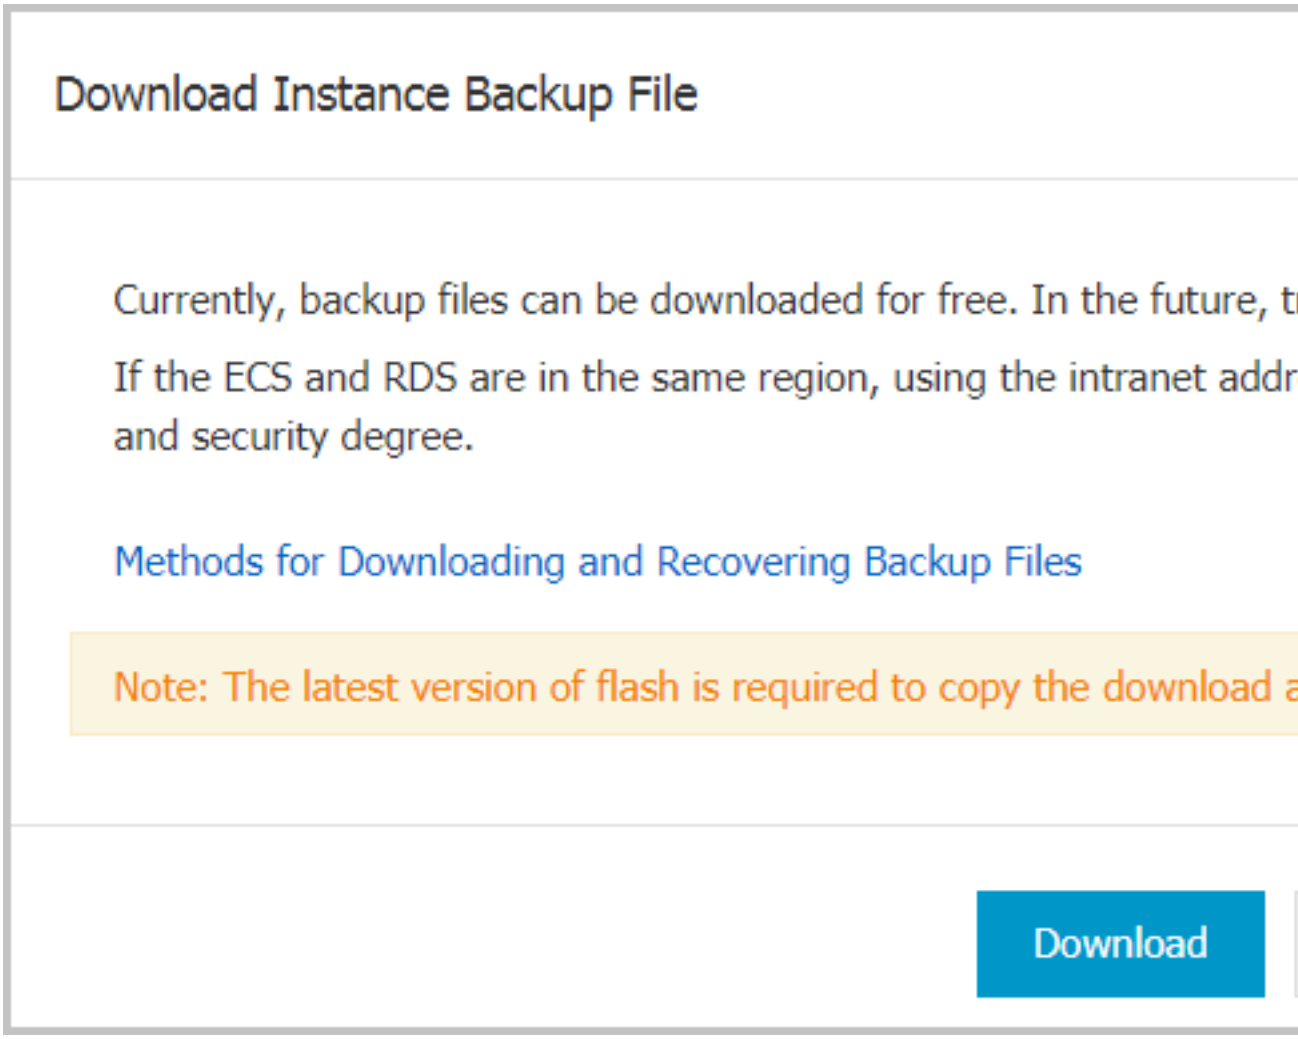

Download method	Description
Download	Directly download the backup file through the Internet.
Copy Intranet Address	If ECS and RDS are in the same region, you can log on to ECS and use the RDS intranet IP address to download the backup file. This method is faster and more secure.
Copy Internet Address	You can copy the Internet IP address and use other tools to download the backup file.

14.4 Logical backup and recovery for PPAS

This document describes the procedure for logical backup and recovery for RDS for PPAS instances.

Procedure

1. Install the PPAS program.

   > **Note:**
   >
   > You must use the PPAS binary system for export. Using the PostgreSQL community binary system leads to an error.
   >
   > - Download Windows client (*Part 1*, *Part 2*)
   > - *Download Linux client (32-bit)*
   > - *Download Linux client (64-bit)*

2. Grant all permissions to a role (to export the data).

   For example, if role A is used to export data but there are two other roles, namely, B and C, in the database, you must run the following commands to grant role A the permissions of role B and role C.

   ```
   -- Use Role B for logon to run the following
   command :
    grant B to A ;
   -- Then use Role A for logon to run the
   following command :
    grant C to A ;
   ```

   In this way, role A has the permission to access all data tables of role B and role C.

3. In the directory where `pg_dump` is located, run the following backup command:

   ```
   ./ pg_dump - h < host > - p < port > - U < user > - f dump .
   sql < dbname >
   ```

4. If recovery is required, you can run the following commands in the directory where `psql` is located:

   ```
   ./ psql - h < host > - p < port > - U < user > - d postgres
   - c " drop database < dbname >"
   ./ psql - h < host > - p < port > - U < user > - d postgres
   - c " create database < dbname >"
   ```

```
./ psql  - h  < host > - p  < port > - U  < user > - f   dump . sql
 - d  < dbname >
```

FAQ

1. The following error occurs when you export data from PPAS:

   ```
   ERROR :  permission   denied   for   relation   product_co
   mponent_ve  rsion
    LOCK   TABLE   sys . product_co  mponent_ve  rsion   IN   ACCESS
   SHARE   MODE
   ```

   Solution: The cause for this error is that you have used the pg_dump program of PG to export data from PPAS. You can use the PPAS binary system to export the data. For PPAS downloading methods, see the preceding procedure.

2. The following error occurs when you export data from PPAS:

   ```
    ERROR :  permission   denied   for   relation  < user   table >
   ```

   Solution: The cause for this error is that the account used for data export has no permission to access the data of other roles. If acceptable, you can grant a role the permissions of other roles and then use this role to export data by running the following command:

   ```
   GRANT   ROLE < other   roles >,< other   roles >  to  < user   for
     pg_dump >
   ```

3. The following error occurs when you use pg_dump.

   ```
    pgdump  - U   xxx  - h   yyy  - p3433  < dbname > - f   my . sql
    pg_dump :  too   many   parameters  ( the   first   one   is  ”- f
   )  in   the   command   line
   ```

   Solution: When running pg_dump on Windows, you must append all other parameters with <dbname>.

4. A parameter error occurs when you use pg_dump.

   Solution: The possible cause is that the specified parameter is incorrect, such as `pg_dump  - Uxxx  - h   yyy` . This parameter is not allowed since a space is needed next to -U (other parameters also follow this style).

15 Recovery

15.1 Restore MySQL data

You can restore data of RDS for MySQL in either of the following ways:

- Restore data to a clone instance. For details, see this article.
- Restore certain databases or tables rather than the entire instance. For details, see *Restore databases or tables for MySQL*.

This article describes how to restore data of the entire instance to a new instance (referred to as a clone instance), verify the data on the clone instance, and transfer the data you need from the clone instance to the original instance.

Note:
The clone instance has the same whitelist, backup settings, and parameter settings as the original instance.

Pricing

The costs are the same as purchasing a new instance. For details, see *Pricing*.

Prerequisites

- The original instance is running properly and not locked.
- The original instance is not undergoing a migration task.
- To restore data to a point in time, ensure that log backup has been enabled.
- To restore data from a backup set, ensure that at least one backup set has been generated.

Attention

- If the data volume is large, the restoration may take a long time.
- If no resource is available when you create a clone instance, try again by choosing a different zone in the same region.

Restore data to a new RDS instance (clone instance)

1. Log on to the *RDS console*.

2. Select the region where the instance is located.

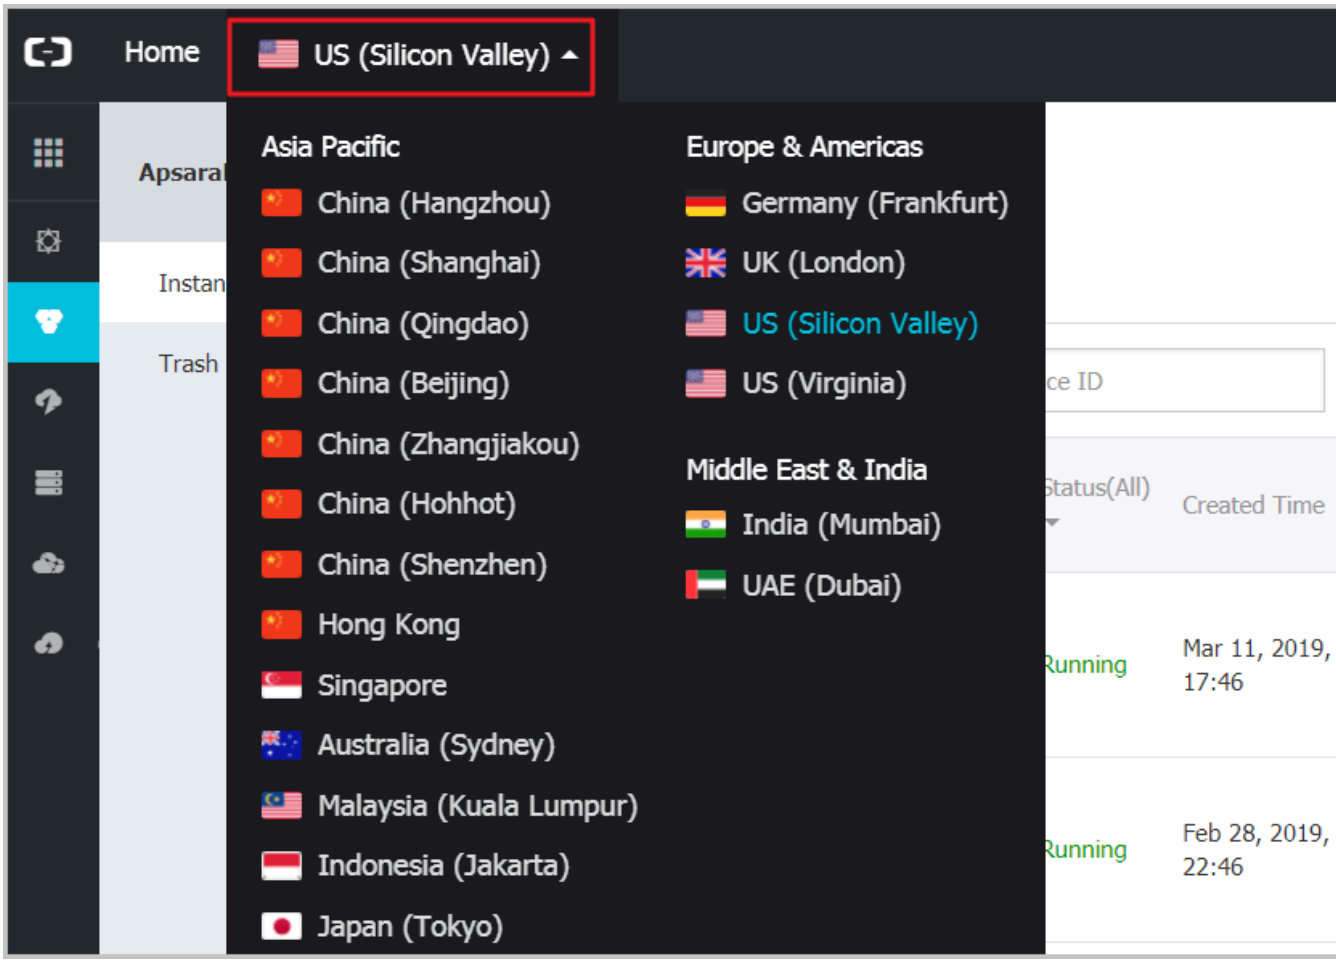

3. Click the instance ID.
4. In the left-side navigation pane, choose Backup and Recovery.
5. In the upper-right corner, click Restore Database.
6. In the displayed window, select a payment method:

   - Pay-As-You-Go: indicates post payment. The system deducts an hourly fee from your account balance every hour. If you plan to use the instance for a short term, this method is cost-effective because you can release the instance after using it.
   - Subscription: indicates prepayment. You need to pay for the instance when creating it. If you plan to use the instance for a month or more, this method is more cost-effective than Pay-As-You-Go. The longer the subscription is, the higher the discount.

Note:

Pay-As-You-Go instances can be changed to Subscription instances, but Subscription instances cannot be changed to Pay-As-You-Go instances.

7. Set the instance parameters.

Parameter	Description
Restore Type	· By Time: You can restore data to any point in time within the log backup retention period. To view or modify the log backup retention period, see *Back up RDS data*. · By Backup ID Note: By Time is displayed only if log backup is enabled.
Edition	· Basic Edition: consists of a single node and separates computing from storage. This edition is cost-effective but is not recommended for production environments. · High-availability: consists of a master node and a slave node. This edition applies to over 80% of application scenarios. For more information, see *Product series overview*.
Zone	A zone is an independent area within a region. Different zones within the same region are basically the same. You can deploy your RDS and ECS instances in the same zone or in different zones. Certain regions allow you deploy a High-availability instance across zones, such as Zone F + Zone G. This indicates that the master and slave nodes of the High-availability instance are in two different zones so that the disaster recovery capability is higher. This does not incur extra costs. Note: The region of the clone instance is the same as that of the original instance.

Parameter	Description
Type	It is recommended that the specifications and storage of the clone instance be equal to higher than those of the original instance; otherwise, the data restoration may take a long time. Each type of specification provides a specific number of CPU cores, memory, maximum number of connections, and maximum IOPS. For details, see *Instance type list*. RDS provides the following instance type families: · General: A general instance has its own memory and I/O resources , and shares CPU and storage resources with other general instances on the same server. · Dedicated: A dedicated instance has it own CPU, memory, storage, and I/O resources. For example, 8 Cores, 32GB is a general instance. 8 Cores, 32GB (Dedicated) is a dedicated instance.
Capacity	The capacity is used for storing data, system files, binlog files, and transaction files.
Network Type	· Classic Network: Traditional network type. · VPC (recommended): VPC is short for Virtual Private Cloud. A VPC is an isolated network and provides higher security and performance than the traditional classic network.

8. Set the duration (for Subscription instances only) and quantity of the instances to be created.
9. Click Buy Now.
10. Review order information, select Product Terms of Service and Service Level Notice and Terms of Use, and complete the payment.

Log on to the clone instance and verify the data

For information about how to log on to an instance, see *Connect to an instance*.

Migrate data to the original instance

After verifying the data on the clone instance, if necessary, you can migrate the data you need from the clone instance to the original instance.

Data migration indicates copying data from one instance (source instance) to another (target instance) and does not affect the source instance.

Attention

DDL operations are not allowed during the migration; otherwise, the migration may fail.

Procedure

1. **Log on to the *DTS console*.**
2. **In the left-side navigation pane, choose Data Migration.**
3. **Click Create Migration Task.**

4. Enter the task name, source database information, and target database information.

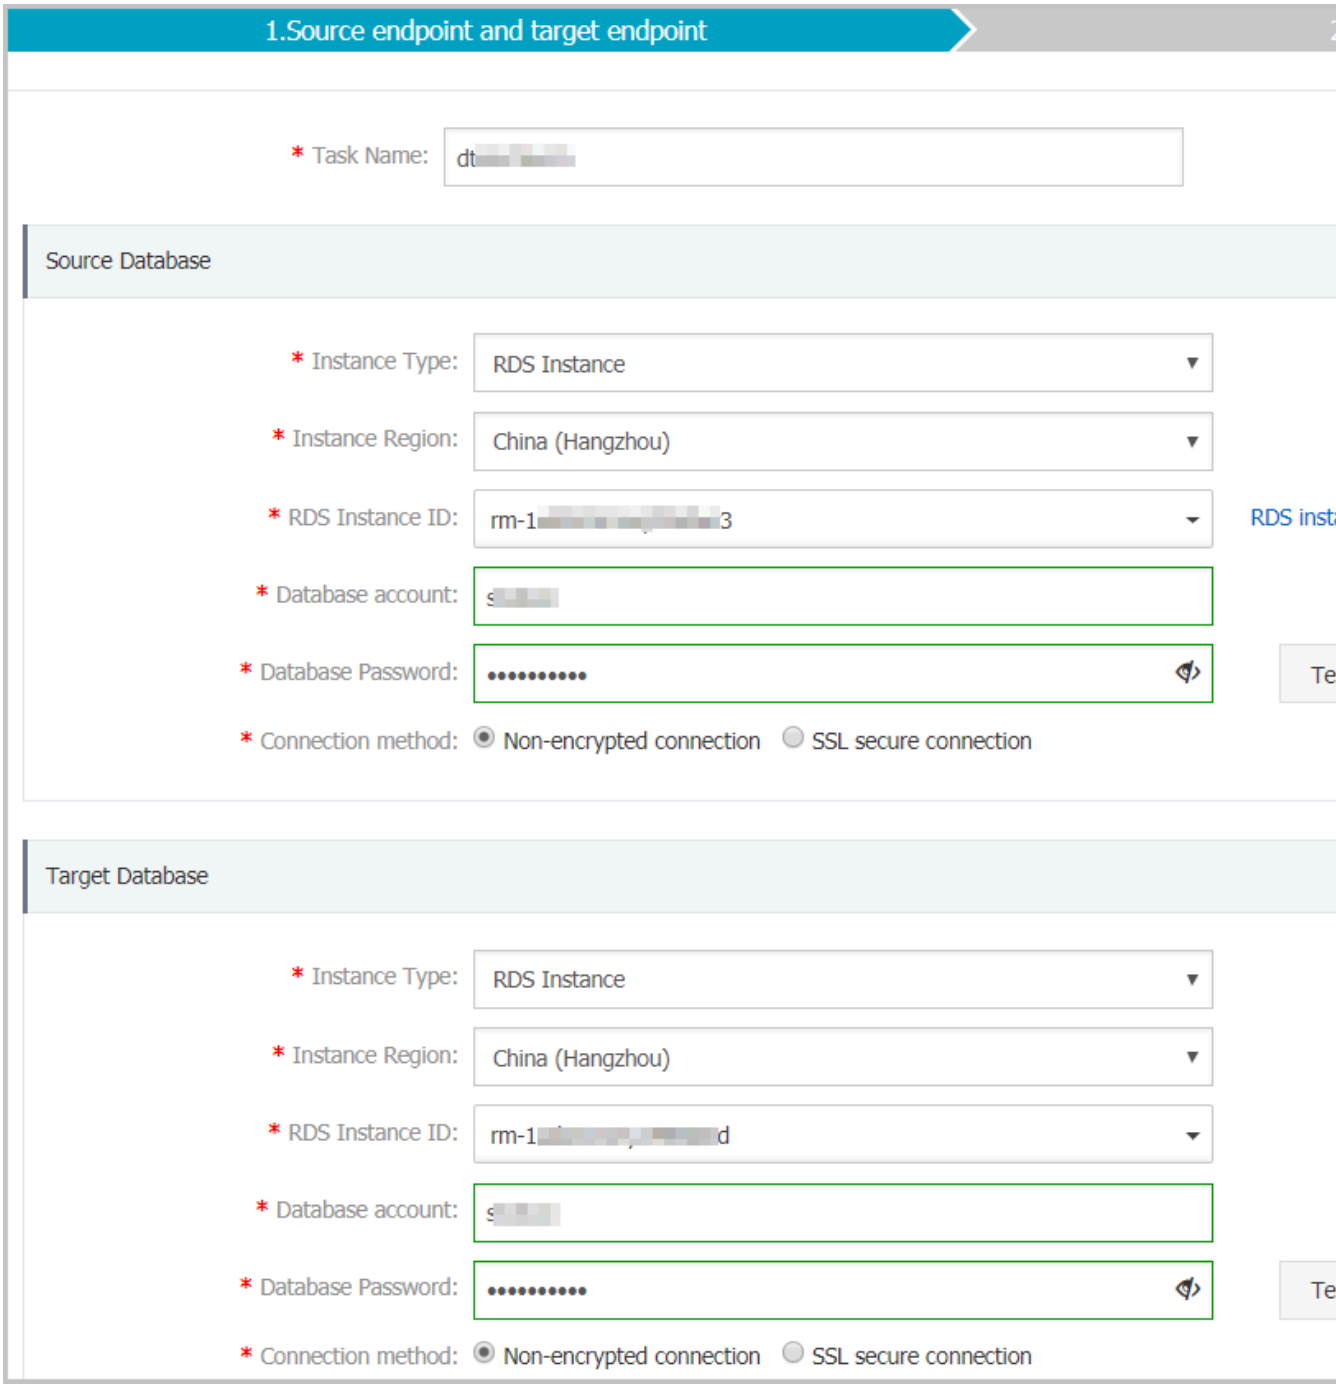

- Task name: A default task name is generated automatically . It is recommended that you set a meaningful name so that the task can be identified easily.
- Source database information:
  - Instance Type: Select RDS Instance.
  - Instance Region: Select the region where the clone instance is located.
  - RDS Instance ID: Select the ID of the clone instance.

> Note:
> This parameter is displayed only if you have selected RDS Instance for Instance Type.

- Database Account: Enter the account name of the clone instance.
- Database Password: Enter the password of the preceding account.
- Connection method: Generally, select Non-encrypted connection. If *SSL encryption* has been enabled for the instance, select SSL secure connection.

> Note:
> This parameter is displayed only if you have selected certain RDS instances.

· Target database information:

- Instance Type: Select RDS Instance.
- Instance Region: Select the region where the original instance is located.
- RDS Instance ID: Select the ID of the original instance.

> Note:
> This parameter is displayed only if you have selected RDS Instance for Instance Type.

- Database Account: Enter the account name of the original instance.
- Database Password: Enter the password of the preceding account.
- Connection method: Generally, select Non-encrypted connection. If *SSL encryption* has been enabled for the instance, select SSL secure connection.

> Note:
> This parameter is displayed only if you have selected certain RDS instances.

5. Click Authorized Whitelist and Enter Into Next Step.
6. Select Migrate object structure and Migrate existing data.
7. In the left pane, select objects and click > to add them to the right.

> Note:
> DTS will perform a data check. If an object in the target instance has the same name as an object to be migrated, the migration fails.

If an object in the target instance has the same name as an object to be migrated, do either of the following:

- In the right pane, place your mouse over an object and click Edit to modify the object name.
- Rename the object in the target instance.

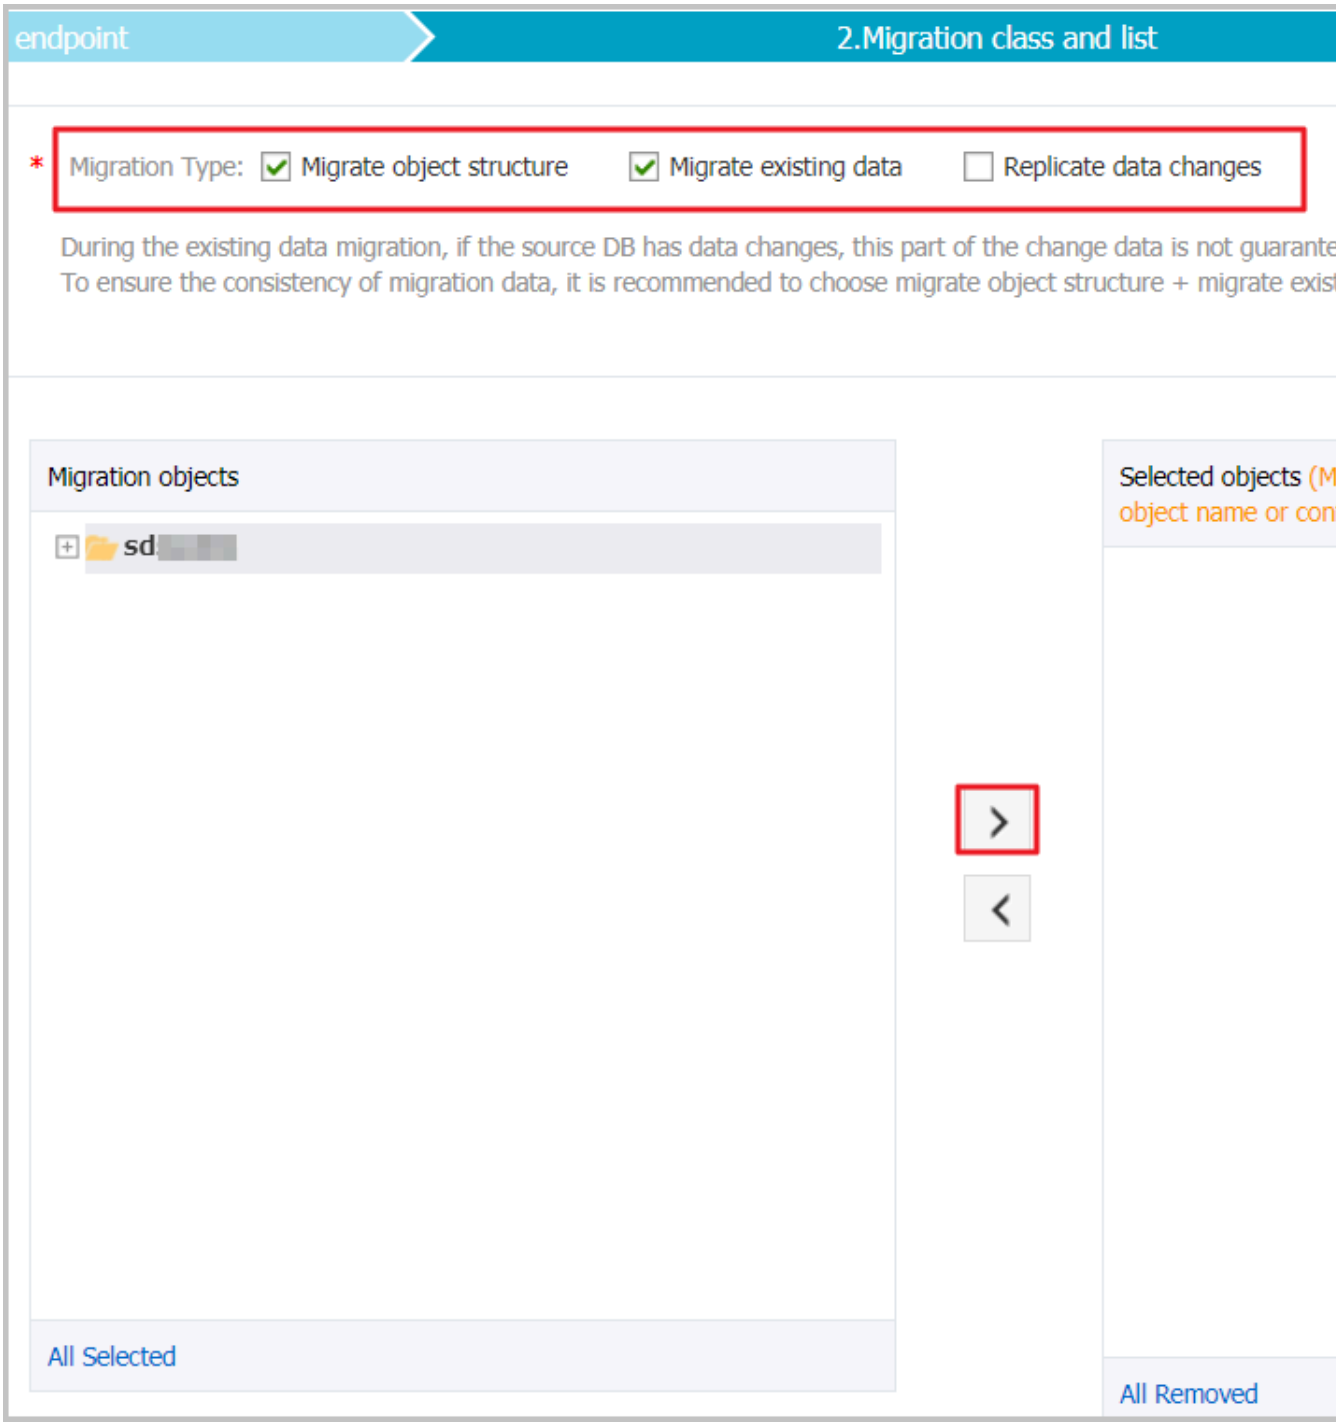

8. Click Pre-check and Start.
   - If the pre-check succeeds, go to step 11.
   - If the pre-check fails, go to step 9.
9. If the pre-check fails, click ⓘ next to the failed item to view details.

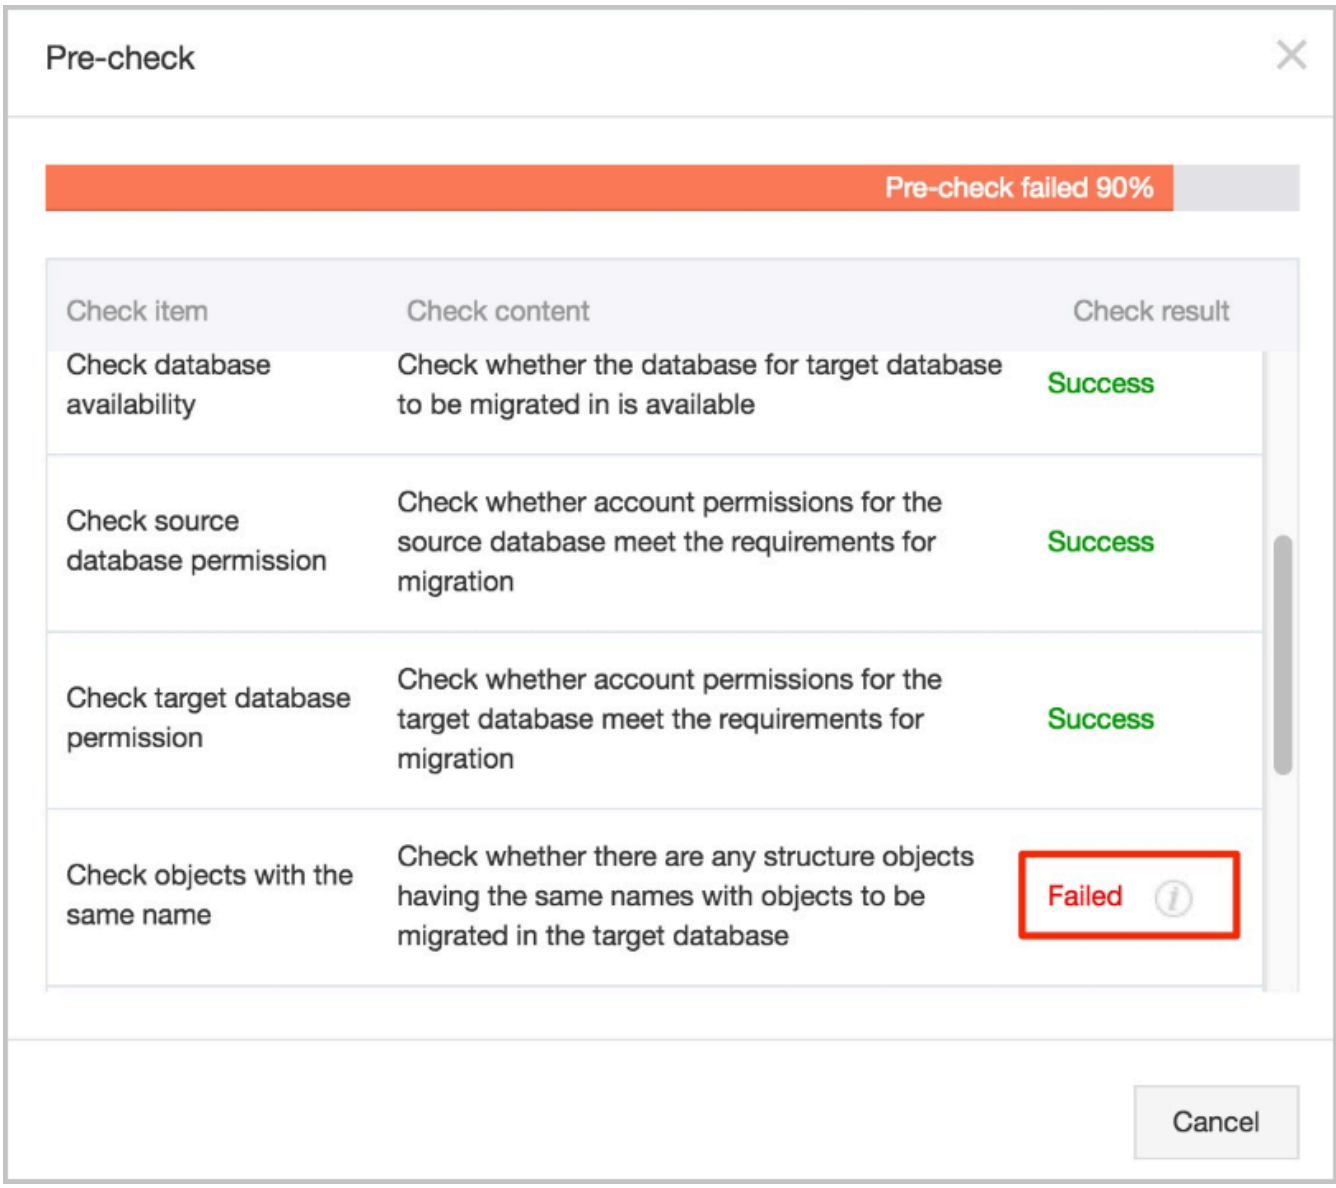

10. After fixing all problems, select the migration task in the migration task list and click Start.

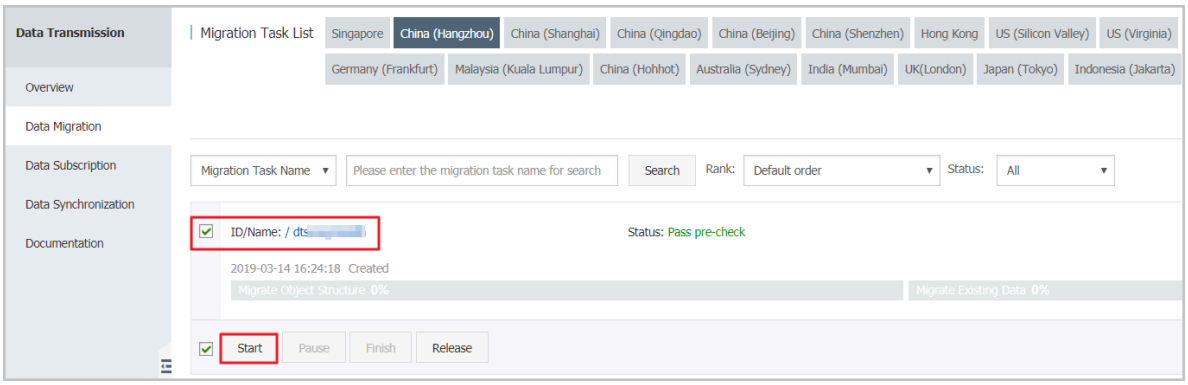

11. If the pre-check succeeds, click Next.
12. On the Confirm Purchase Configuration dialog box, confirm the configuration, select Service Terms of Data Transmission (Pay-As-You-Go), and click Buy and Start Now.

15.2 Restore MySQL databases or tables

In RDS for MySQL 5.6 High-Availability Edition, you can restore only certain databases or tables rather than the entire instance.

Prerequisites

- The instance type is RDS for MySQL 5.6 High-Availability Edition.
- The region is Singapore. If your instance is in other regions, please *submit a ticket* to apply for activating the function.
- The instance is running properly and not locked.
- To restore data from a backup set, the instance must have at least one backup set.
- To restore data to a point in time, make sure that the log backup function is enabled.

Precautions

After this function is activated, the backup file format is changed from TAR to XBSTREAM, so the backup files occupy a little more space. Pay attention to the *backup file size* because the excess space that exceeds the free quota will incur costs. You can adjust the backup frequency if needed.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the instance is located.
3. Click the instance ID.
4. In the left-side navigation pane, choose Backup and Recovery.
5. In the upper right corner, click Restore Databases/Tables.

> **Note:**
> If this button is not displayed, see Prerequisites in this topic.

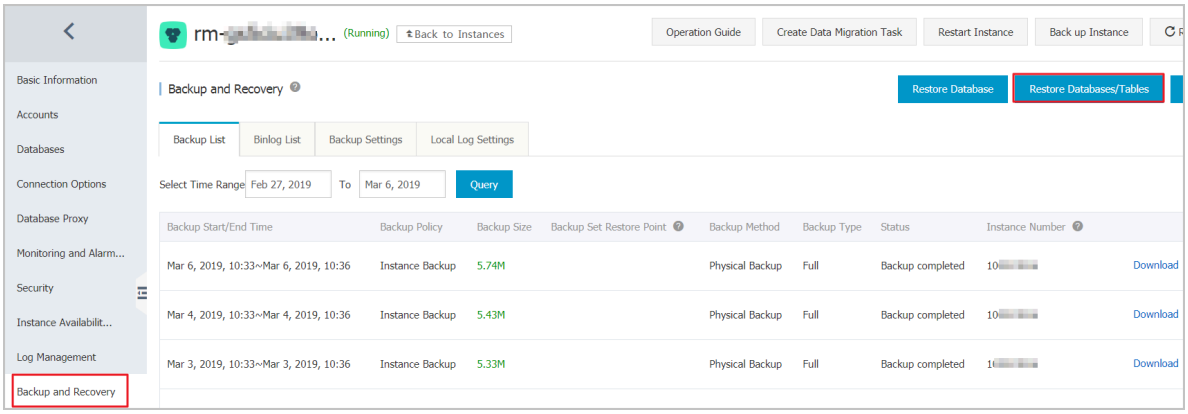

6. Set the following parameters.

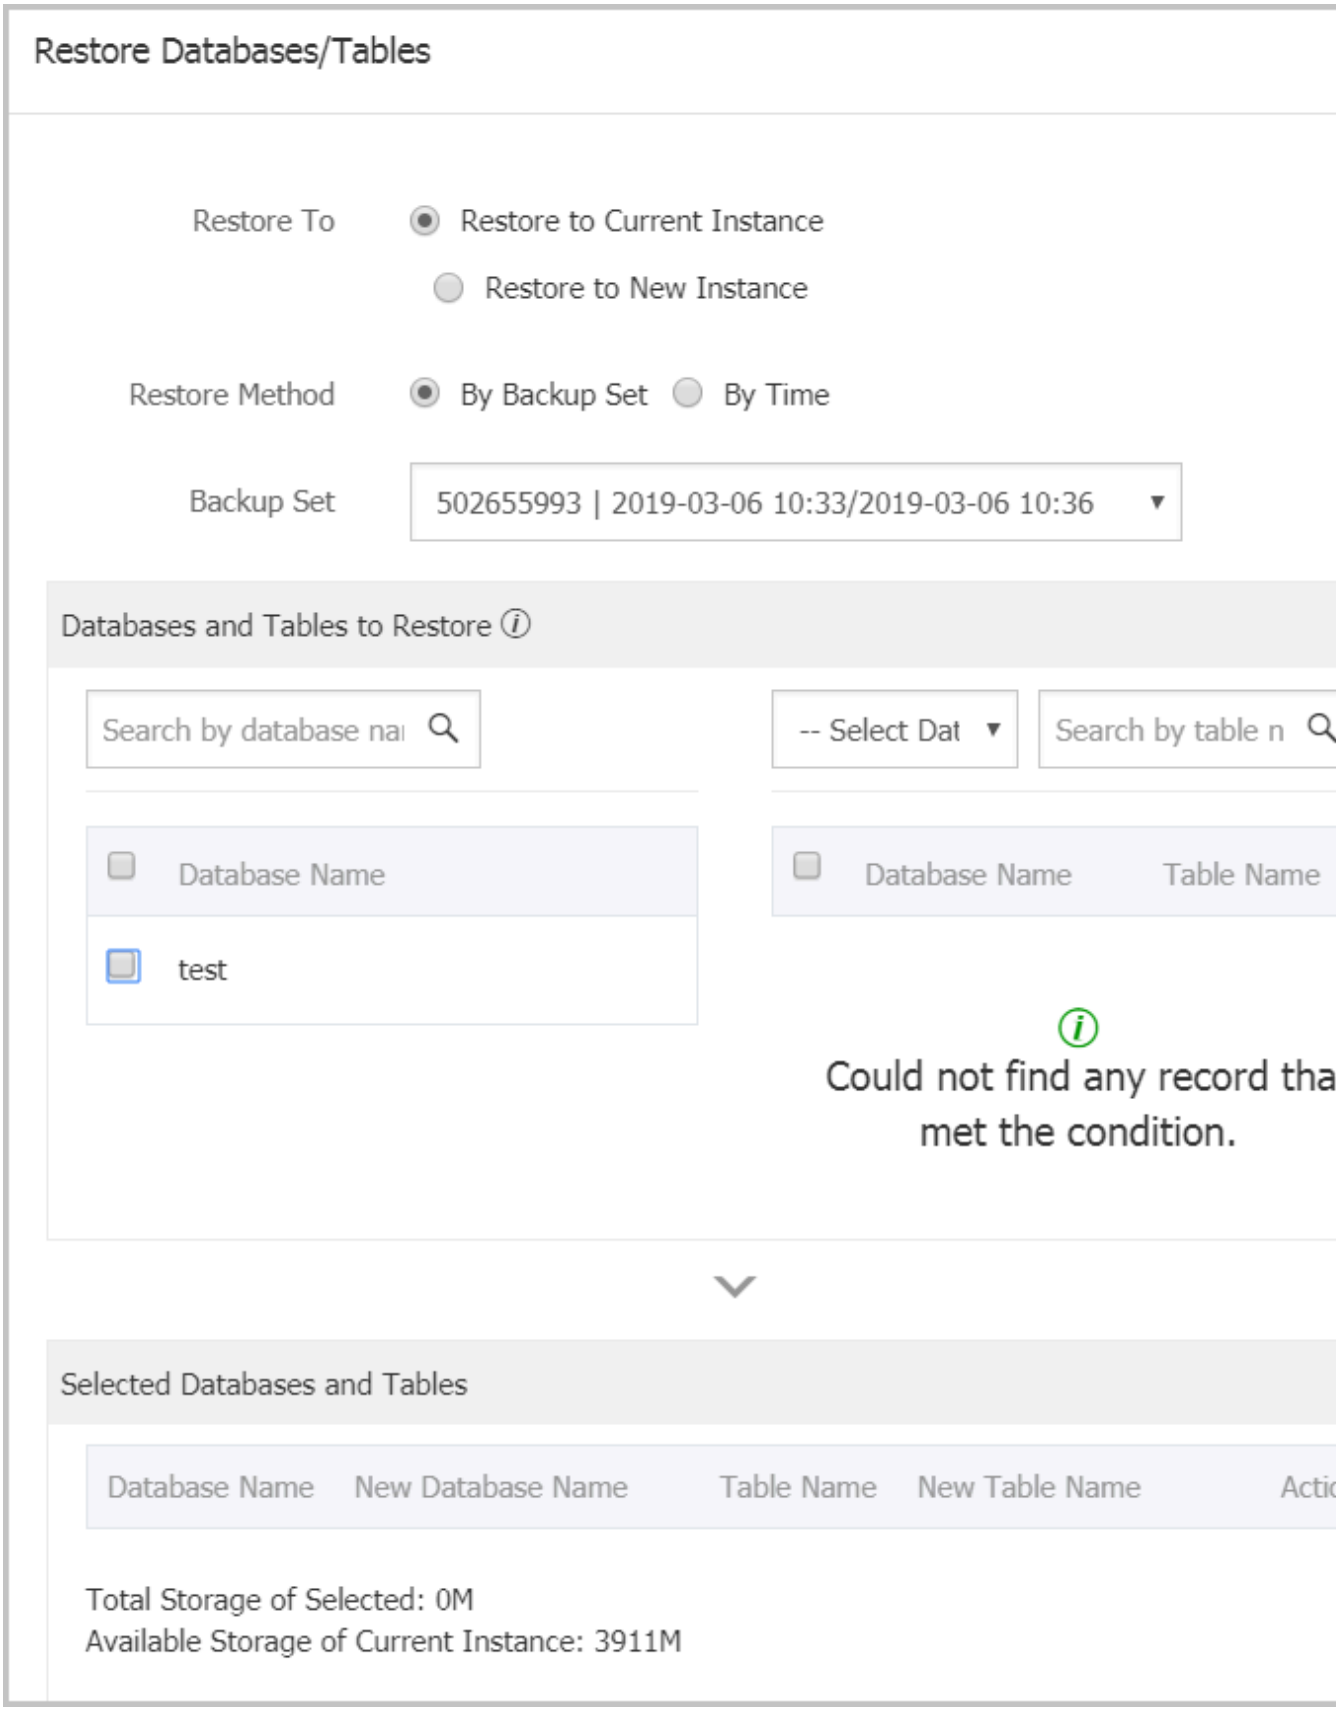

Parameter	Description
Restore To	· Restore to Current Instance: If you select this option, ensure that the instance is not undergoing a migration process. · Restore to New Instance
Restore Method	· By Backup Set. · By Time. This parameter is displayed only if the log backup function is enabled.
Backup Set	Select a backup set. Note: This parameter is displayed only if Restore Method is By Backup Set.
Restore Time	Select a point in time. You can restore data to any time within the log retention period. To view or modify the log retention period, see *Back up RDS data*. Note: This parameter is displayed only if Restore Method is By Time.
Databases and Tables to Restore	Select the databases or tables to restore.
Selected Databases and Tables	Selected databases and tables are displayed here. · If needed, you can set the database and table names that are used after the restoration. · This area also displays the total size of the selected databases and tables and the available storage of the current instance.

7. Click OK. The restoration starts.

Note:
If you chose to restore a new instance, the instance purchase page is displayed.
After you complete the payment, the restoration starts.

15.3 Restore SQL Server Data

You can restore data of RDS for SQL Server in any of the following ways.

· *Restore to an existing RDS instance*

- *Restore to new RDS instance*
- *Restore to a temporary RDS instance*

Attention

If the data volume is large, the restoration may take a long time.

Restore data to an existing RDS instance

You can restore all or part of the databases in your instance to an existing RDS instance.

Applicable scope

This method applies to RDS for SQL Server 2016 or 2012 instances.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the instance is located.

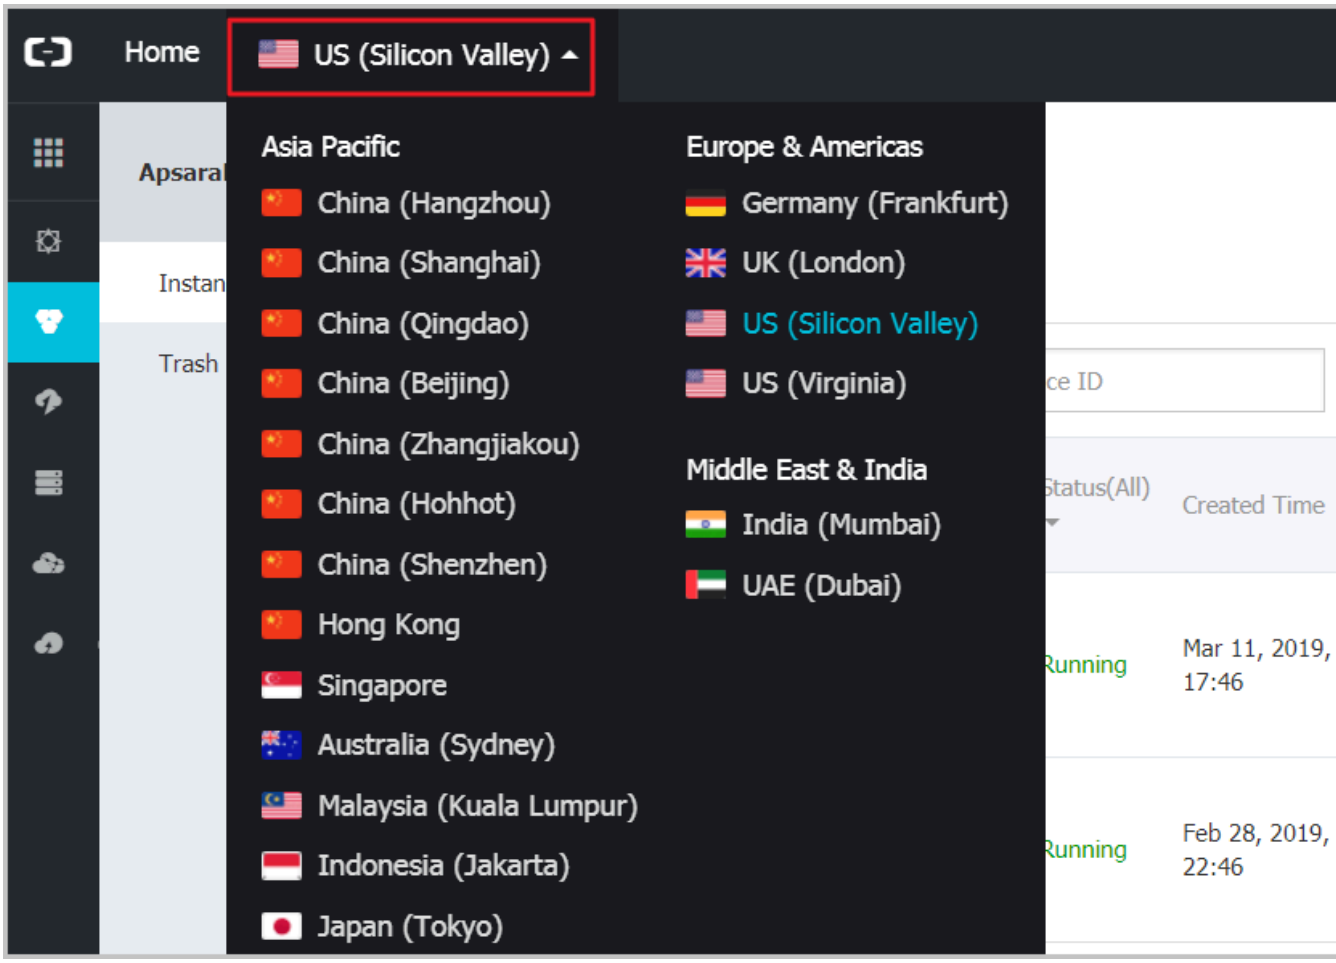

3. Click the instance ID.
4. In the left-side navigation pane, choose Backup and Recovery.

5. In the upper-right corner of the page, click Restore.
6. (This step is for high-availability series only.) SelectRestore to Existing Instance and click OK.

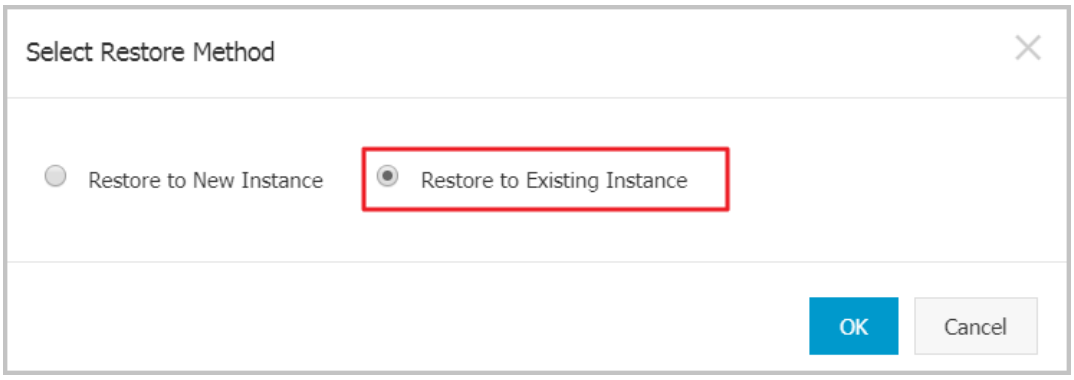

7. Set the following parameters, and then clickOK.

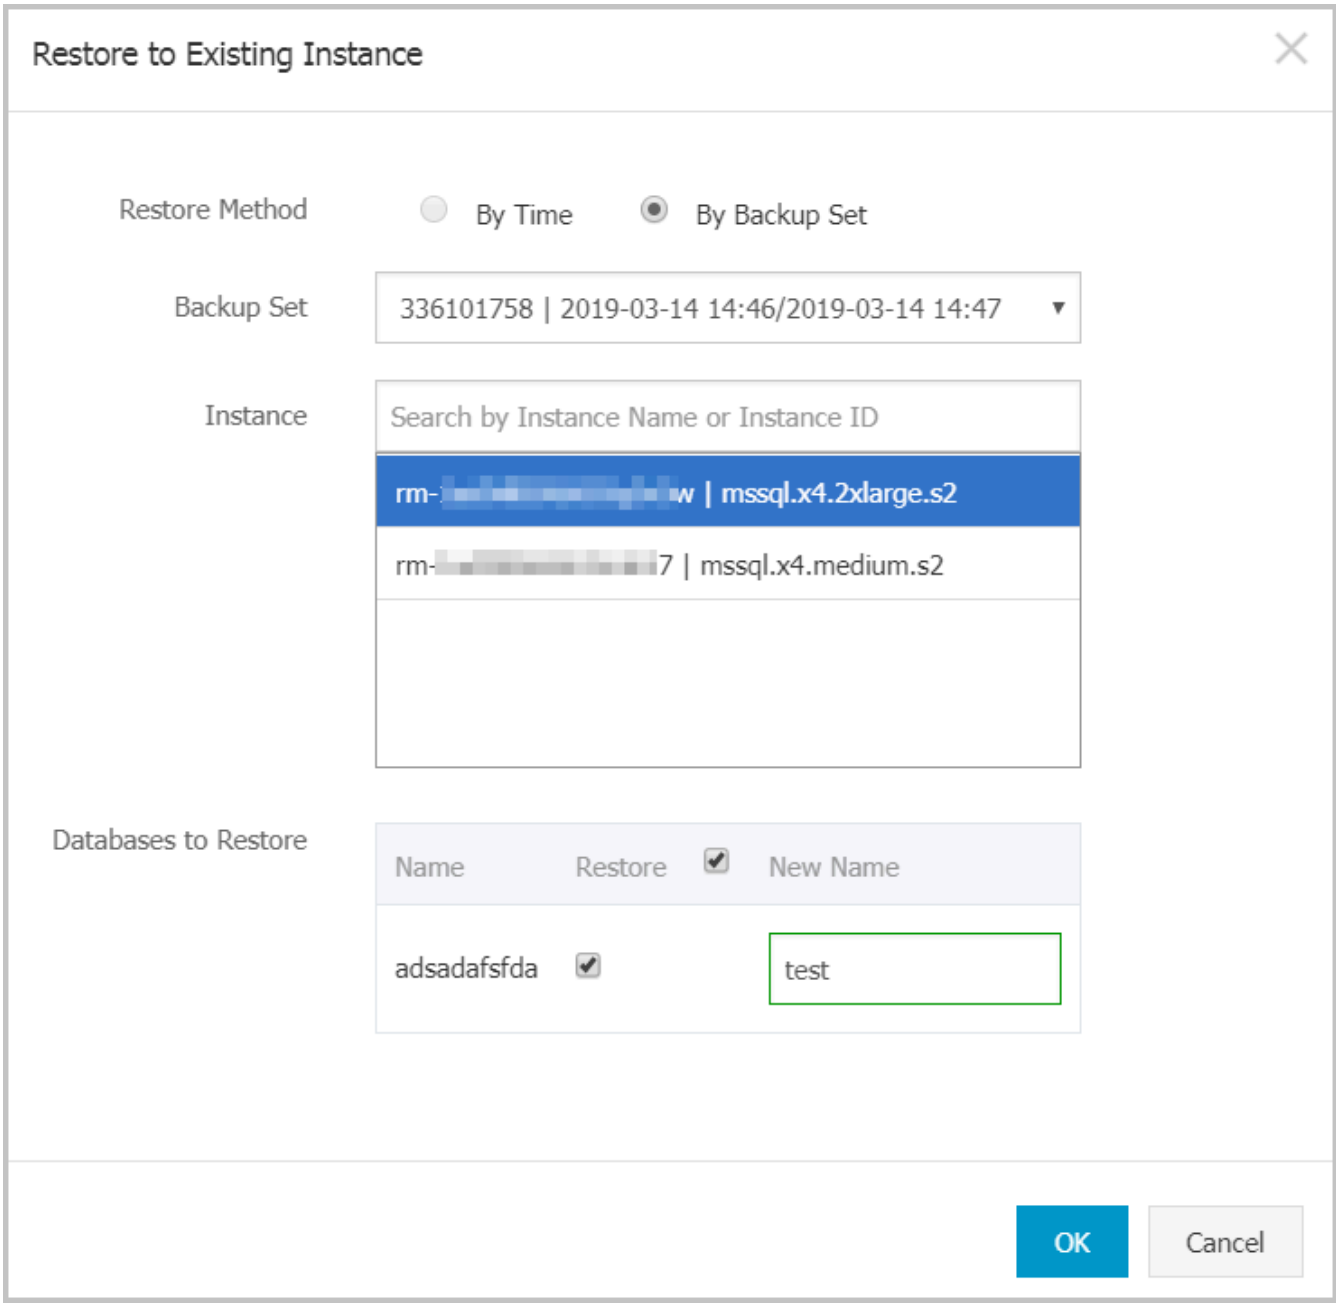

Note:

If the existing instance already has a database that has the same name as the database to be restored, you need to modify `New Name` .

Parameter	Description
Restore Method	To restore data to an existing instance, select By Backup Set.

Parameter	Description
Backup Set	Select the backup set to restore. By default, the system displays all full backup sets under the current instance.
Instance	Select the instance to which the backup set will be restored. By default, the system displays the current instance and all instances that belong to the current Alibaba account and current region and have the same database version as the current instance. Note: If many instances are displayed, you can use the search box.
Databases to restore	a. Select the database to restore. All databases in the backup set are displayed and selected by default. · To restore data of the entire instance, retain the default selection (All databases are selected). · To restore certain databases, select only these databases. b. Set the database names that are displayed after the databases are restored. By default, the database names in the backup set are used. Note: The database names cannot be the same as the existing database names in the target instance.

Restore to a new RDS instance

This function is also called "clone instance", used to restore the historical backup of the instance to a new instance. You can restore data by time or backup set. When restoring by backup set, you can restore all or part of the databases in the backup set.

Pricing

The costs are the same as purchasing a new instance. For details, see *Pricing*.

Applicable scope

This method applies to the following instances:

· SQL Server 2017 Cluster series
· SQL Server 2012/2016 Enterprise Edition High-Availability series
· SQL Server 2012/2016 Standard Edition High-Availability series

Procedure

1. Log on to the *RDS console*.
2. Select the region where the instance is located.

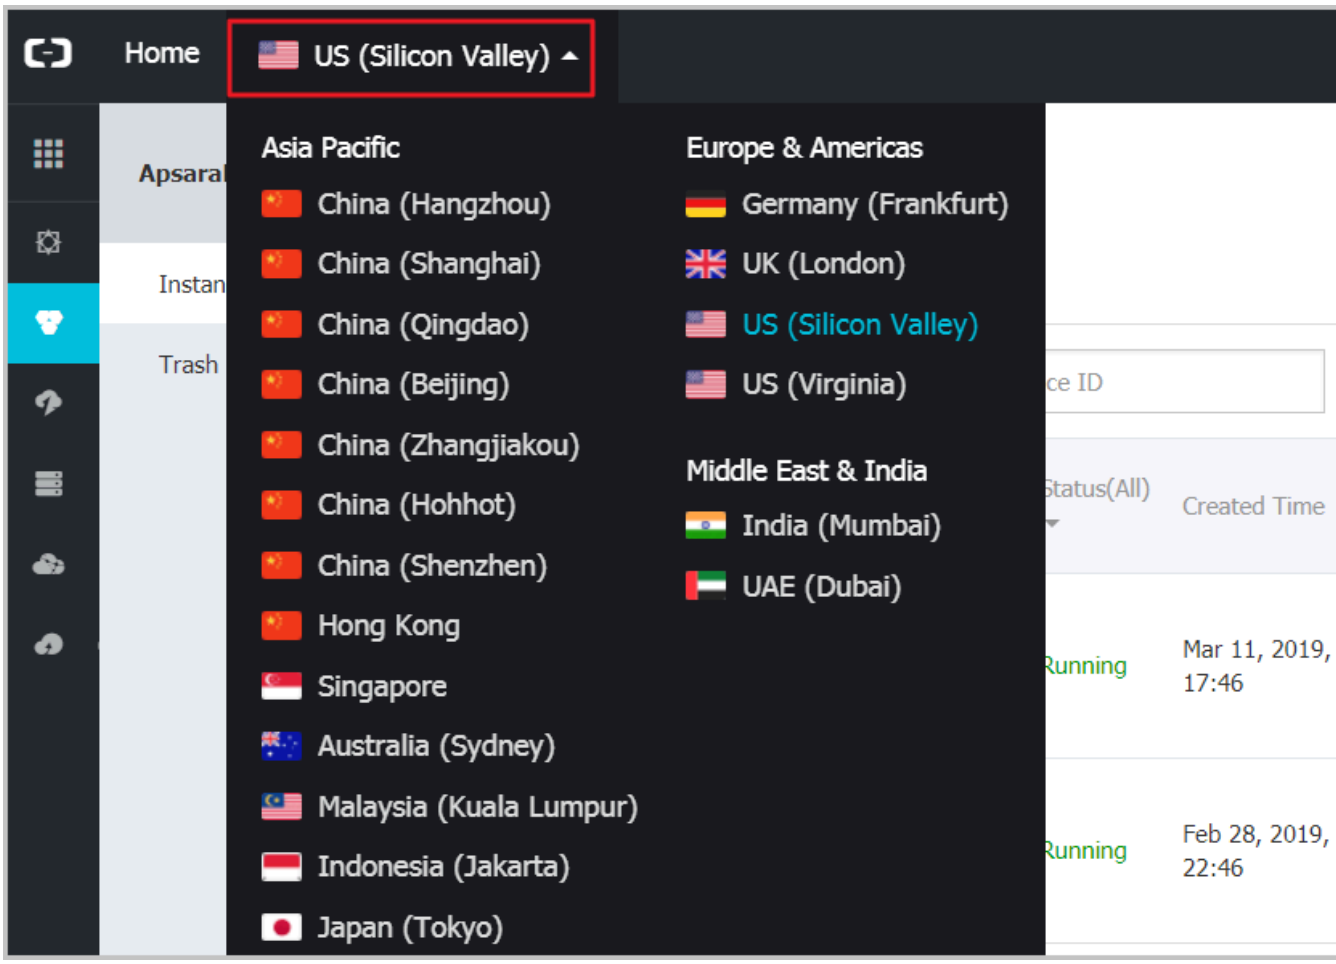

3. Click the instance ID.
4. In the left-side navigation pane, choose Backup and Recovery.
5. In the upper-right corner of the page, click Restore.
6. Select Restore to New Instance and click OK.
7. In the displayed window, select a payment method:
   - Pay-As-You-Go: indicates post payment. The system deducts an hourly fee from your account balance every hour. If you plan to use the instance for a short term, this method is cost-effective because you can release the instance after using it.
   - Subscription: indicates prepayment. You need to pay for the instance when creating it. If you plan to use the instance for a month or more, this method

is more cost-effective than Pay-As-You-Go. The longer the subscription is, the higher the discount.

> Note:
> Pay-As-You-Go instances can be changed to Subscription instances, but Subscription instances cannot be changed to Pay-As-You-Go instances.

8. Set the instance parameters.

Parameter	Description
Restore Type	· By Time: You can restore data to any point in time within the log backup retention period. To view or modify the log backup retention period, see *Back up RDS data*. · By Backup ID Note: By Time is displayed only if log backup is enabled.
Database	· All: Restore all databases in the backup set. · Part: Restore part of the databases in the backup set.
Edition	· High-availability: consists of a master node and a slave node. This edition applies to over 80% of application scenarios. · AlwaysOn (Cluster) Edition: provides one master node, one slave node, and up to seven read-only nodes that horizontally scale read capabilities. For more information, see *Cluster Edition (AlwaysOn Edition)*.
Zone	A zone is an independent area within a region. Different zones within the same region are basically the same. You can deploy your RDS and ECS instances in the same zone or in different zones. Note: The region of the clone instance is the same as that of the original instance.

Parameter	Description
Type	It is recommended that the specifications and storage of the clone instance be equal to higher than those of the original instance; otherwise, the data restoration may take a long time. Each type of specification provides a specific number of CPU cores, memory, maximum number of connections, and maximum IOPS. For details, see *Instance type list*. RDS provides the following instance type families: · General: A general instance has its own memory and I/O resources , and shares CPU and storage resources with other general instances on the same server. · Dedicated: A dedicated instance has it own CPU, memory, storage, and I/O resources. For example, 8 Cores, 32GB is a general instance. 8 Cores, 32GB (Dedicated) is a dedicated instance.
Capacity	The capacity is used for storage data, system files, and transaction files.
Network Type	· Classic Network: Traditional network type. · VPC (recommended): VPC is short for Virtual Private Cloud. A VPC is an isolated network and provides higher security and performance than the traditional classic network.

9. Click Buy Now.
10. Review order information, select Product Terms of Service and Service Level Notice and Terms of Use, and complete the payment.

Restore data to a temporary instance

This method applies to the following instances:

· SQL Server 2012 Enterprise Edition Basic series
· SQL Server 2012/2016 Web Edition Basic series
· SQL Server 2008 R2

For detailed operations, see*Recover data to a temporary instance (RDS for SQL Server)*.

15.4 Restore PostgreSQL or PPAS data

RDS for PostgreSQL/PPAS allows you to restore data by time or backup set. The restoration process is as follows:

- Restore data of an RDS instance to a new RDS instance (referred to as a clone instance).
- Verify the data on the clone instance.
- Migrate the data you need from the clone instance to the original instance.

Note:

- The clone instance has the same whitelist, backup settings, and parameter settings as the original instance.
- RDS for PostgreSQL and PPAS does not allow you to restore data of an instance directly to the instance itself to overwrite the existing data.

Pricing

The costs are the same as purchasing a new instance. For details, see *Pricing*.

Prerequisites

- The original instance is running properly and not locked.
- The original instance is not undergoing a migration task.
- To restore data to a point in time, ensure that log backup has been enabled.
- To restore data from a backup set, ensure that at least one backup set has been generated.

Attention

- If the data volume is large, the restoration may take a long time.
- If no resource is available when you create a clone instance, try again by choosing a different zone in the same region.

Restore data to a new RDS instance (clone instance)

1. Log on to the *RDS console*.

2. Select the region where the instance is located.

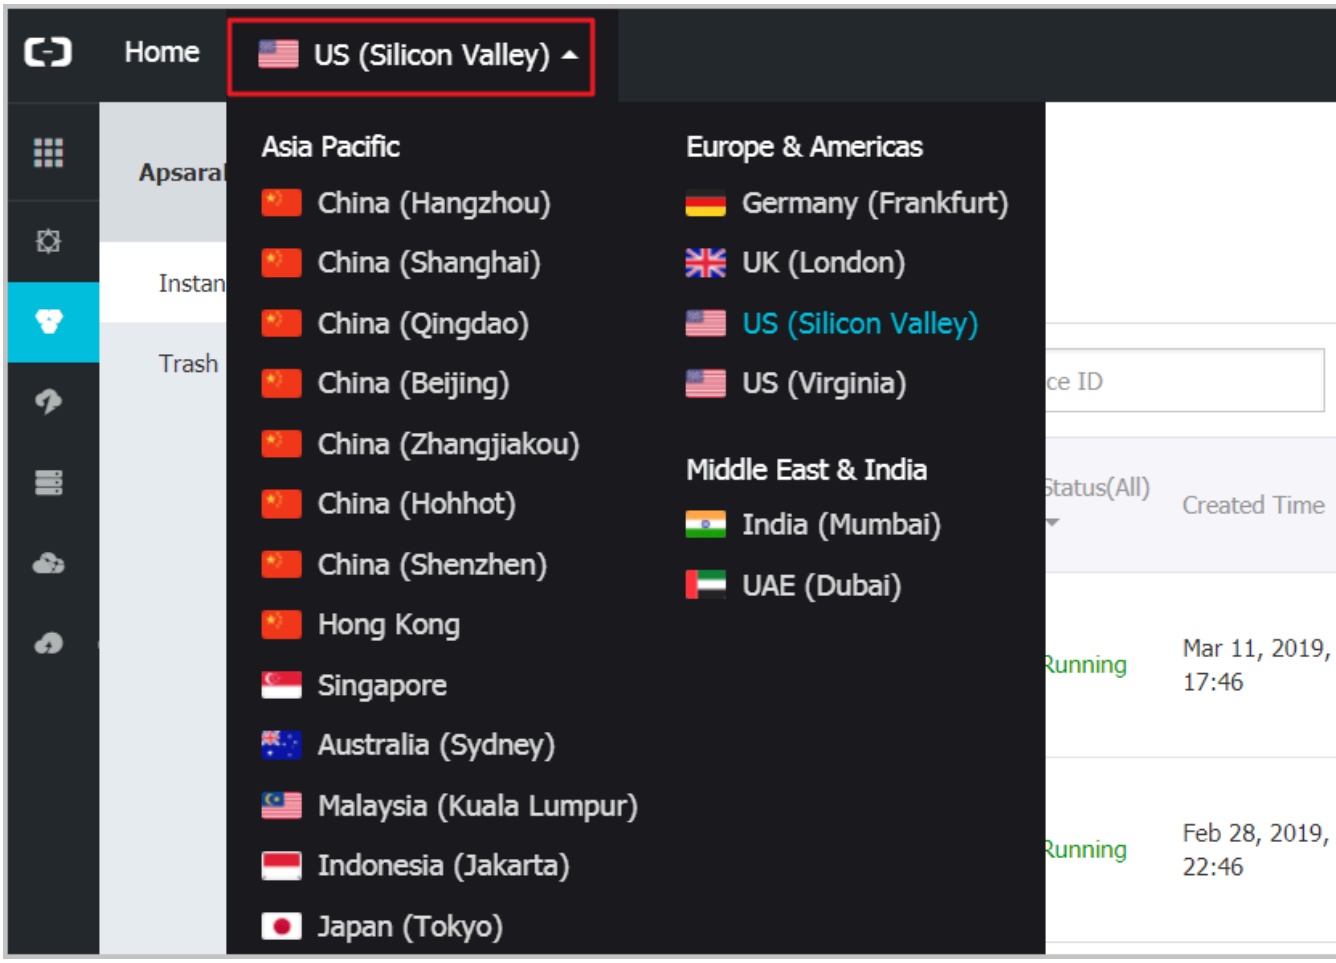

3. Click the instance ID.
4. In the left-side navigation pane, choose Backup and Recovery.
5. In the upper-right corner, click Restore Database.
6. In the displayed window, select a payment method:

   - Pay-As-You-Go: indicates post payment. The system deducts an hourly fee from your account balance every hour. If you plan to use the instance for a short term, this method is cost-effective because you can release the instance after using it.
   - Subscription: indicates prepayment. You need to pay for the instance when creating it. If you plan to use the instance for a month or more, this method is more cost-effective than Pay-As-You-Go. The longer the subscription is, the higher the discount.

Note:

Pay-As-You-Go instances can be changed to Subscription instances, but Subscription instances cannot be changed to Pay-As-You-Go instances.

7. Set the instance parameters.

Parameter	Description
Restore Type	· By Time: You can restore data to any point in time within the log backup retention period. To view or modify the log backup retention period, see *Back up RDS data*. · By Backup ID Note: By Time is displayed only if log backup is enabled.
Edition	· Basic Edition: consists of a single node and separates computing from storage. This edition is cost-effective but is not recommended for production environments. · High-availability: consists of a master node and a slave node. This edition applies to over 80% of application scenarios. For more information, see *Product series overview*.
Zone	A zone is an independent area within a region. Different zones within the same region are basically the same. You can deploy your RDS and ECS instances in the same zone or in different zones. Certain regions allow you deploy a High-availability instance across zones, such as Zone F + Zone G. This indicates that the master and slave nodes of the High-availability instance are in two different zones so that the disaster recovery capability is higher. This does not incur extra costs. Note: The region of the clone instance is the same as that of the original instance.

Parameter	Description
Type	It is recommended that the specifications and storage of the clone instance be equal to higher than those of the original instance; otherwise, the data restoration may take a long time. Each type of specification provides a specific number of CPU cores, memory, maximum number of connections, and maximum IOPS. For details, see *Instance type list*. RDS provides the following instance type families: · General: A general instance has its own memory and I/O resources , and shares CPU and storage resources with other general instances on the same server. · Dedicated: A dedicated instance has it own CPU, memory, storage, and I/O resources. For example, 8 Cores, 32GB is a general instance. 8 Cores, 32GB (Dedicated) is a dedicated instance.
Capacity	The capacity is used for storing data and system files.
Network Type	· Classic Network: Traditional network type. · VPC (recommended): VPC is short for Virtual Private Cloud. A VPC is an isolated network and provides higher security and performance than the traditional classic network.

8. Set the duration (for Subscription instances only) and quantity of the instances to be created.
9. Click Buy Now.
10. Review order information, select Product Terms of Service and Service Level Notice and Terms of Use, and complete the payment.

Log on to the clone instance and verify the data

For information about how to log on to an instance, see *Connect to an instance*.

Migrate data to the original instance

After verifying the data on the clone instance, if necessary, you can migrate the data you need from the clone instance to the original instance.

Data migration indicates copying data from one instance (source instance) to another (target instance) and does not affect the source instance.

Attention

DDL operations are not allowed during the migration; otherwise, the migration may fail.

Procedure

1. Log on to the *DTS console*.
2. In the left-side navigation pane, choose Data Migration.
3. Click Create Migration Task.
4. Enter the task name, source database information, and target database information.
   - Task name: A default task name is generated automatically . It is recommended that you set a meaningful name so that the task can be identified easily.
   - Source database information:
     - Instance Type: Select RDS Instance.
     - Instance Region: Select the region where the clone instance is located.
     - RDS Instance ID: Select the ID of the clone instance.

       Note:
       This parameter is displayed only if you have selected RDS Instance for Instance Type.
     - Database Account: Enter the account name of the clone instance.
     - Database Password: Enter the password of the preceding account.
     - Connection method: Generally, select Non-encrypted connection. If *SSL encryption* has been enabled for the instance, select SSL secure connection.

       Note:
       This parameter is displayed only if you have selected certain RDS instances.
   - Target database information:
     - Instance Type: Select RDS Instance.
     - Instance Region: Select the region where the original instance is located.
     - RDS Instance ID: Select the ID of the original instance.

       Note:

This parameter is displayed only if you have selected RDS Instance for Instance Type.

- Database Account: Enter the account name of the original instance.
- Database Password: Enter the password of the preceding account.
- Connection method: Generally, select Non-encrypted connection. If *SSL encryption* has been enabled for the instance, select SSL secure connection.

Note:

This parameter is displayed only if you have selected certain RDS instances.

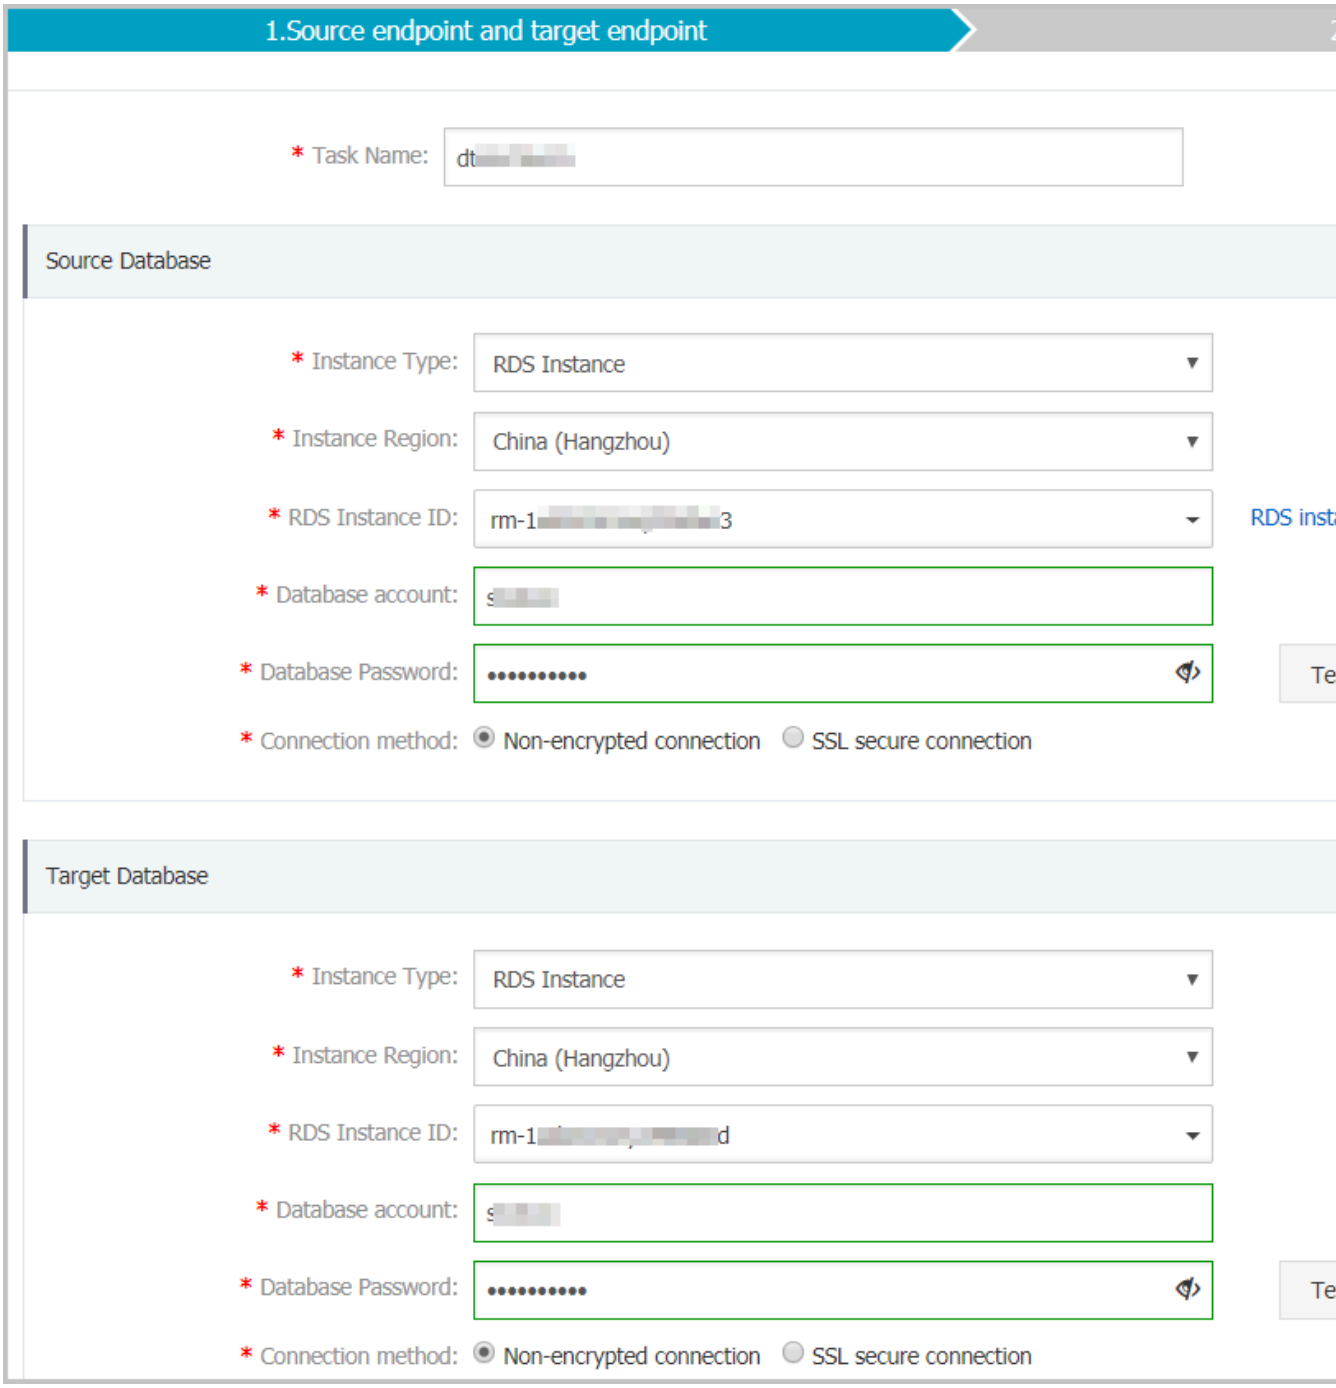

5. Click Authorized Whitelist and Enter Into Next Step.
6. Select Migrate object structure and Migrate existing data.
7. In the left pane, select objects and click > to add them to the right.

Note:

DTS will perform a data check. If an object in the target instance has the same name as an object to be migrated, the migration fails.

If an object in the target instance has the same name as an object to be migrated, do either of the following:

- In the right pane, place your mouse over an object and click Edit to modify the object name.
- Rename the object in the target instance.

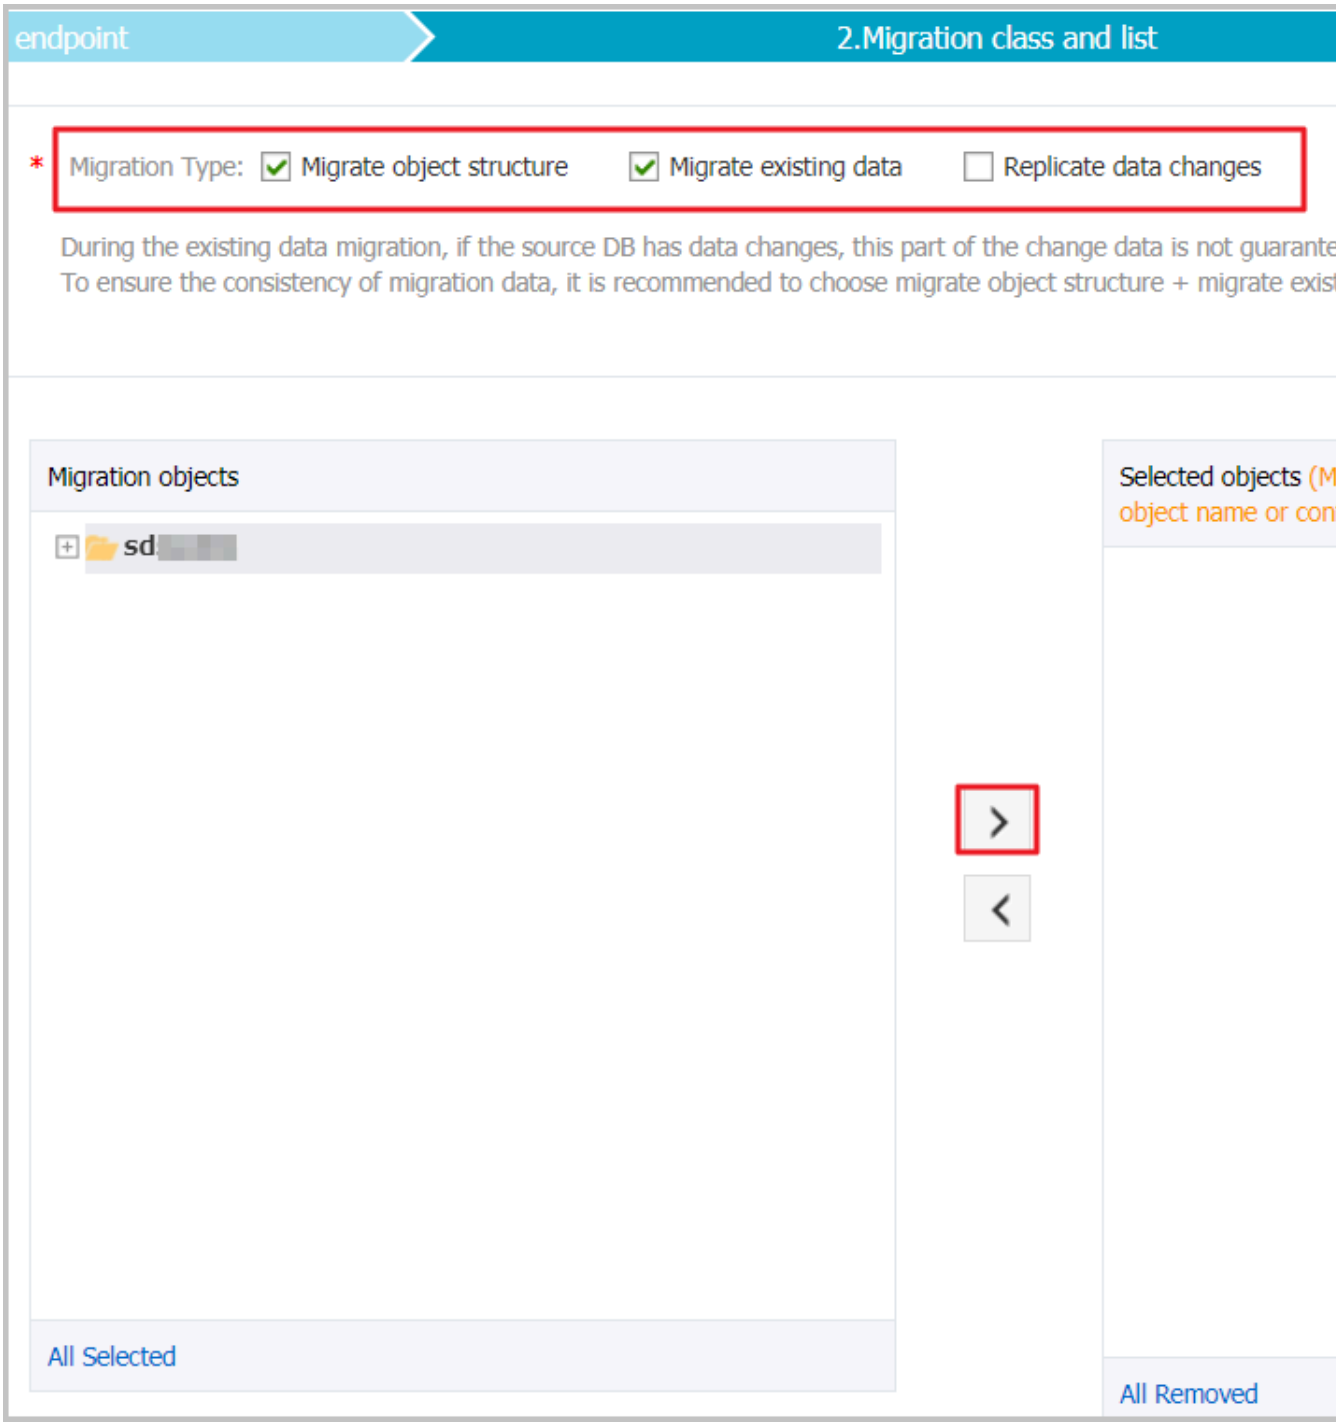

8. Click Pre-check and Start.
   - If the pre-check succeeds, go to step 11.
   - If the pre-check fails, go to step 9.
9. If the pre-check fails, click ⓘ next to the failed item to view details.

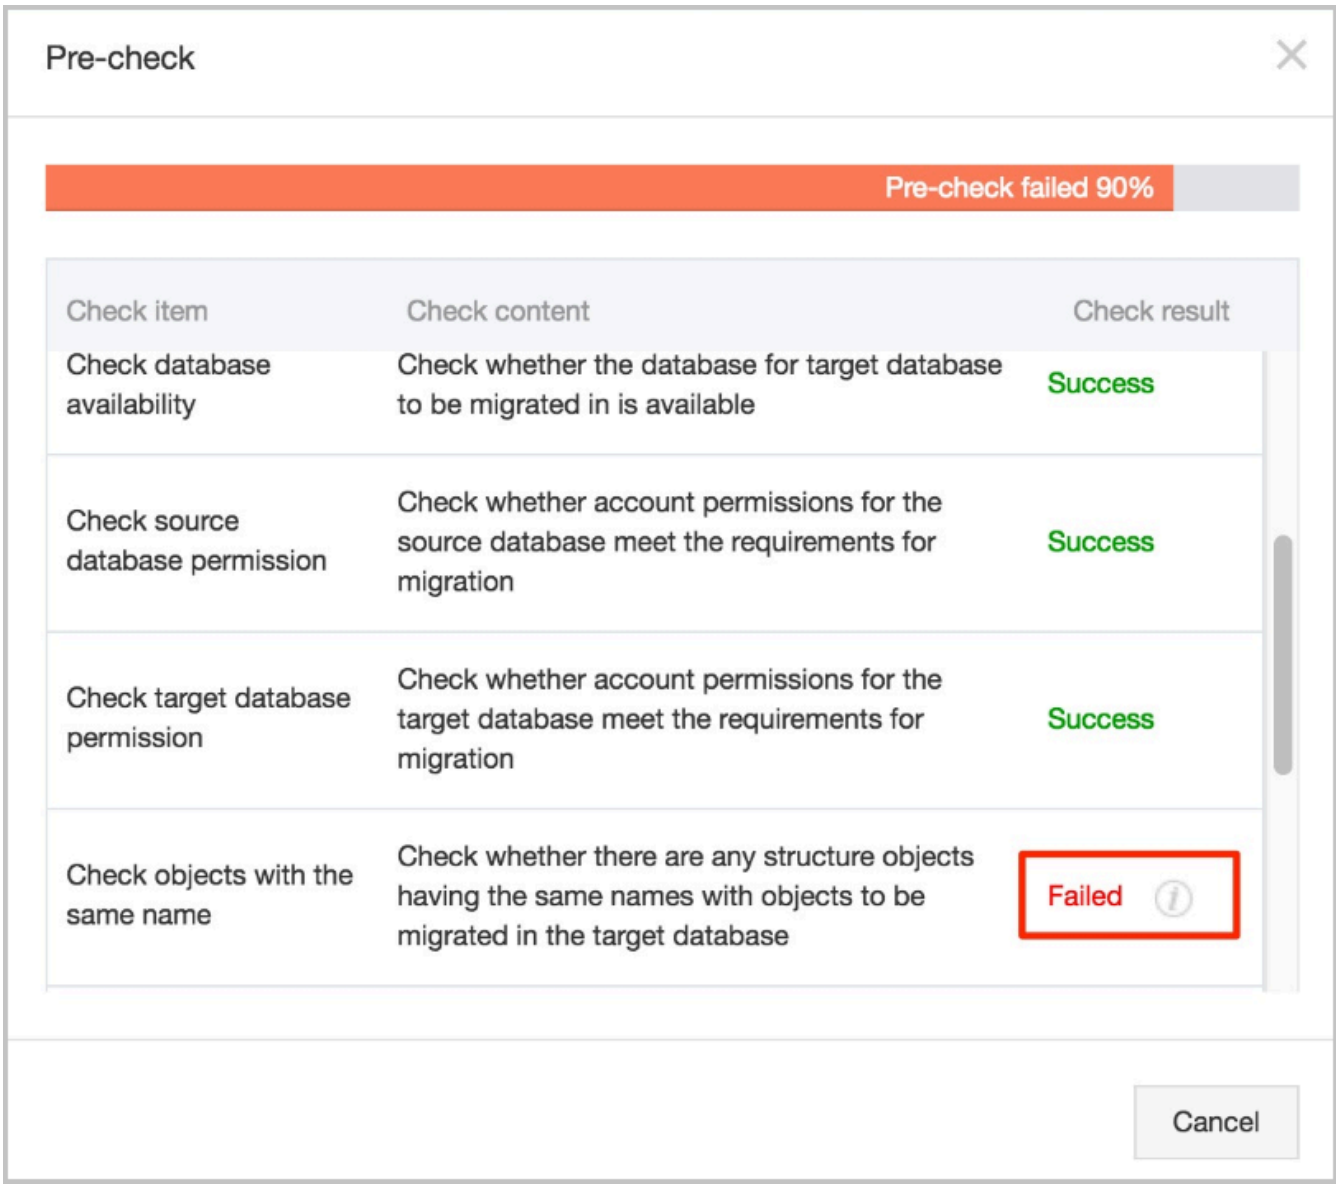

10. After fixing all problems, select the migration task in the migration task list and click Start.

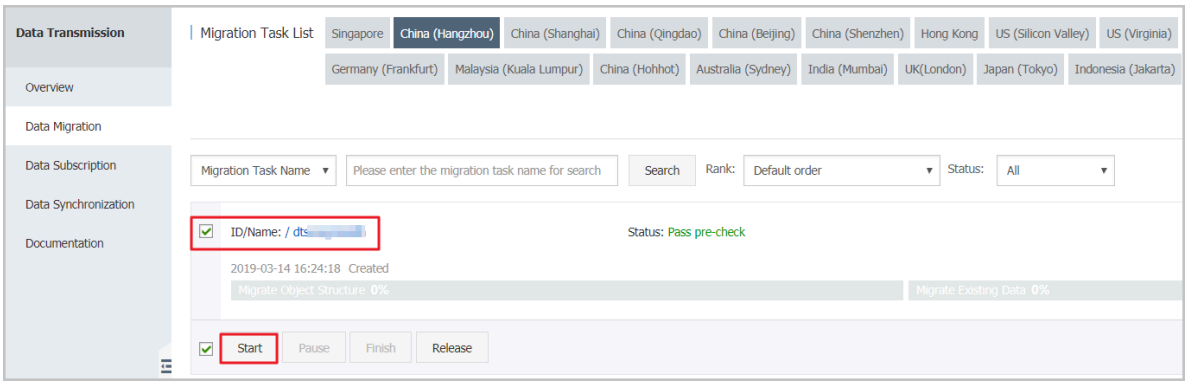

11. If the pre-check succeeds, click Next.
12. On the Confirm Purchase Configuration dialog box, confirm the configuration, select Service Terms of Data Transmission (Pay-As-You-Go), and click Buy and Start Now.

15.5 Restore MariaDB data

You can restore data of RDS for MariaDB TX as follows:

- Restore data of an RDS instance to a new RDS instance (referred to as a clone instance).
- Verify the data on the clone instance.
- Migrate the data you need from the clone instance to the original instance.

> Note:
>
> - The clone instance has the same whitelist, backup settings, and parameter settings as the original instance.
> - RDS for MariaDB TX does not allow you to restore data of an instance directly to the instance itself to overwrite the existing data.

Pricing

The costs are the same as purchasing a new instance. For details, see *Pricing*.

Prerequisites

- The original instance is running properly and not locked.
- The original instance is not undergoing a migration task.
- To restore data to a point in time, ensure that log backup has been enabled.

- To restore data from a backup set, ensure that at least one backup set has been generated.

Attention

- If the data volume is large, the restoration may take a long time.
- If no resource is available when you create a clone instance, try again by choosing a different zone in the same region.

Restore data to a new RDS instance (clone instance)

1. Log on to the *RDS console*.
2. Select the region where the instance is located.

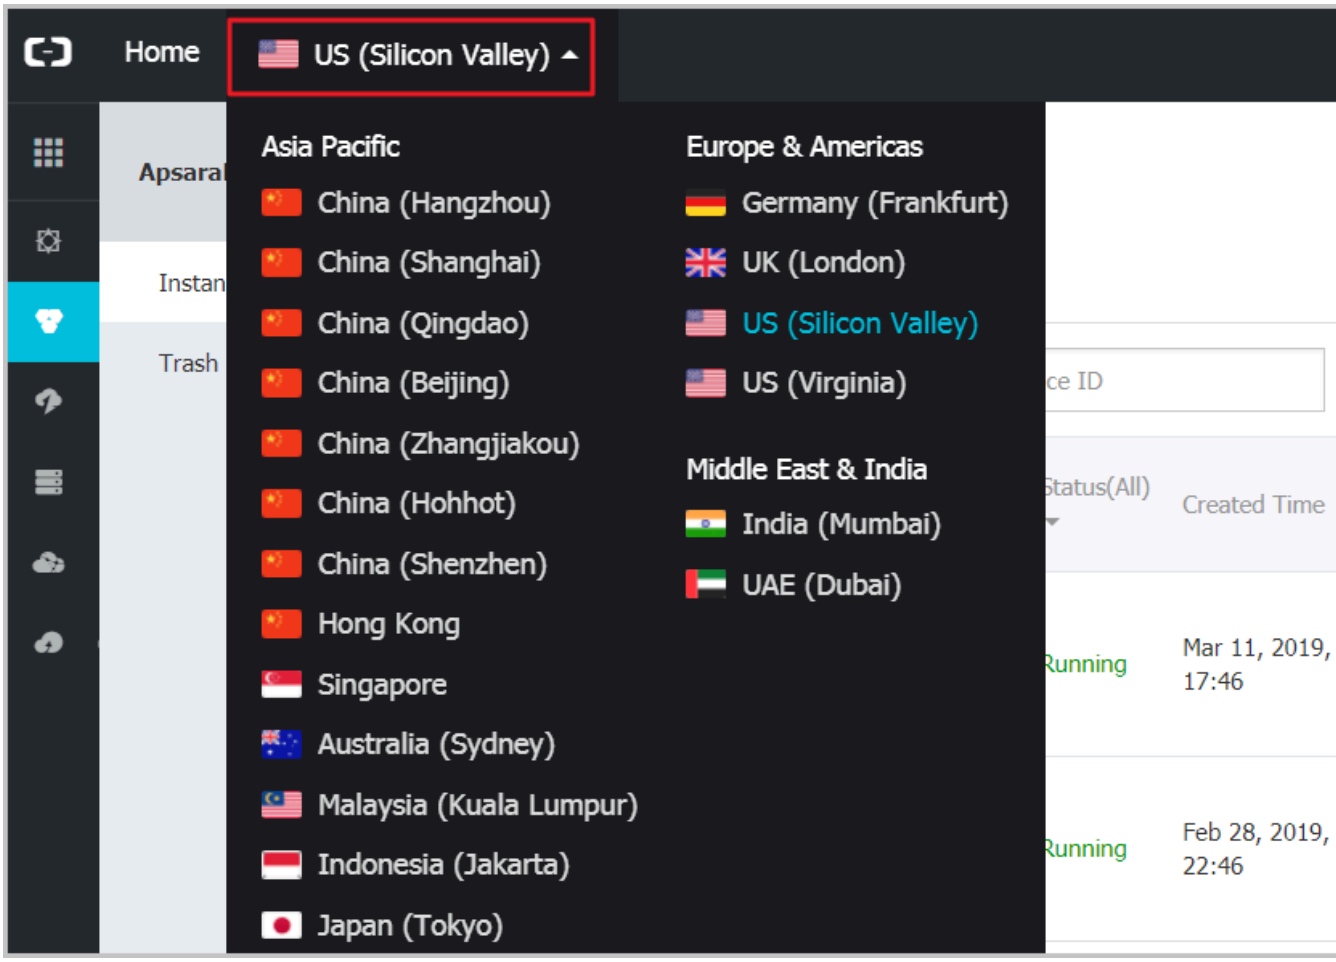

3. Click the instance ID.
4. In the left-side navigation pane, choose Backup and Recovery.
5. In the upper-right corner, click Restore Database.
6. In the displayed window, select a payment method:
   - Pay-As-You-Go: indicates post payment. The system deducts an hourly fee from your account balance every hour. If you plan to use the instance for a short

term, this method is cost-effective because you can release the instance after using it.

- Subscription: indicates prepayment. You need to pay for the instance when creating it. If you plan to use the instance for a month or more, this method is more cost-effective than Pay-As-You-Go. The longer the subscription is, the higher the discount.

> Note:
> Pay-As-You-Go instances can be changed to Subscription instances, but Subscription instances cannot be changed to Pay-As-You-Go instances.

7. Set the instance parameters.

Parameter	Description
Restore Type	· By Time: You can restore data to any point in time within the log backup retention period. To view or modify the log backup retention period, see *Back up RDS data*. · By Backup ID Note: By Time is displayed only if log backup is enabled.
Edition	RDS for MariaDB TX currently supports the High-availability Edition, which consists of a master node and a slave node. This edition applies to over 80% of application scenarios. For more information, see *Product series overview*.
Zone	A zone is an independent area within a region. Different zones within the same region are basically the same. You can deploy your RDS and ECS instances in the same zone or in different zones. Certain regions allow you deploy a High-availability instance across zones, such as Zone F + Zone G. This indicates that the master and slave nodes of the High-availability instance are in two different zones so that the disaster recovery capability is higher. This does not incur extra costs. Note: The region of the clone instance is the same as that of the original instance.

Parameter	Description
Type	It is recommended that the specifications and storage of the clone instance be equal to higher than those of the original instance; otherwise, the data restoration may take a long time. Each type of specification provides a specific number of CPU cores, memory, maximum number of connections, and maximum IOPS. For details, see *Instance type list*. RDS provides the following instance type families: · General: A general instance has its own memory and I/O resources , and shares CPU and storage resources with other general instances on the same server. · Dedicated: A dedicated instance has it own CPU, memory, storage, and I/O resources. For example, 8 Cores, 32GB is a general instance. 8 Cores, 32GB (Dedicated) is a dedicated instance.
Capacity	The capacity is used for storing data, system files, binlog files, and transaction files.
Network Type	RDS for MariaDB TX supports the VPC (short for Virtual Private Cloud). A VPC is an isolated network and provides higher security and performance than the traditional classic network.

8. Set the duration (for Subscription instances only) and quantity of the instances to be created.
9. Click Buy Now.
10. Review order information, select Product Terms of Service and Service Level Notice and Terms of Use, and complete the payment.

Log on to the clone instance and verify the data

For information about how to log on to an instance, see *Connect to an instance*.

Migrate data to the original instance

After verifying the data on the clone instance, if necessary, you can migrate the data you need from the clone instance to the original instance.

Data migration indicates copying data from one instance (source instance) to another (target instance) and does not affect the source instance.

Attention

DDL operations are not allowed during the migration; otherwise, the migration may fail.

Procedure

1. Log on to the *DTS console*.
2. In the left-side navigation pane, choose Data Migration.
3. Click Create Migration Task.
4. Enter the task name, source database information, and target database information.
   - Task name: A default task name is generated automatically . It is recommended that you set a meaningful name so that the task can be identified easily.
   - Source database information:
     - Instance Type: Select RDS Instance.
     - Instance Region: Select the region where the clone instance is located.
     - RDS Instance ID: Select the ID of the clone instance.

       Note:
       This parameter is displayed only if you have selected RDS Instance for Instance Type.
     - Database Account: Enter the account name of the clone instance.
     - Database Password: Enter the password of the preceding account.
     - Connection method: Generally, select Non-encrypted connection. If *SSL encryption* has been enabled for the instance, select SSL secure connection.

       Note:
       This parameter is displayed only if you have selected certain RDS instances.
   - Target database information:
     - Instance Type: Select RDS Instance.
     - Instance Region: Select the region where the original instance is located.
     - RDS Instance ID: Select the ID of the original instance.

       Note:

This parameter is displayed only if you have selected RDS Instance for Instance Type.

- Database Account: Enter the account name of the original instance.
- Database Password: Enter the password of the preceding account.
- Connection method: Generally, select Non-encrypted connection. If *SSL encryption* has been enabled for the instance, select SSL secure connection.

Note:

This parameter is displayed only if you have selected certain RDS instances.

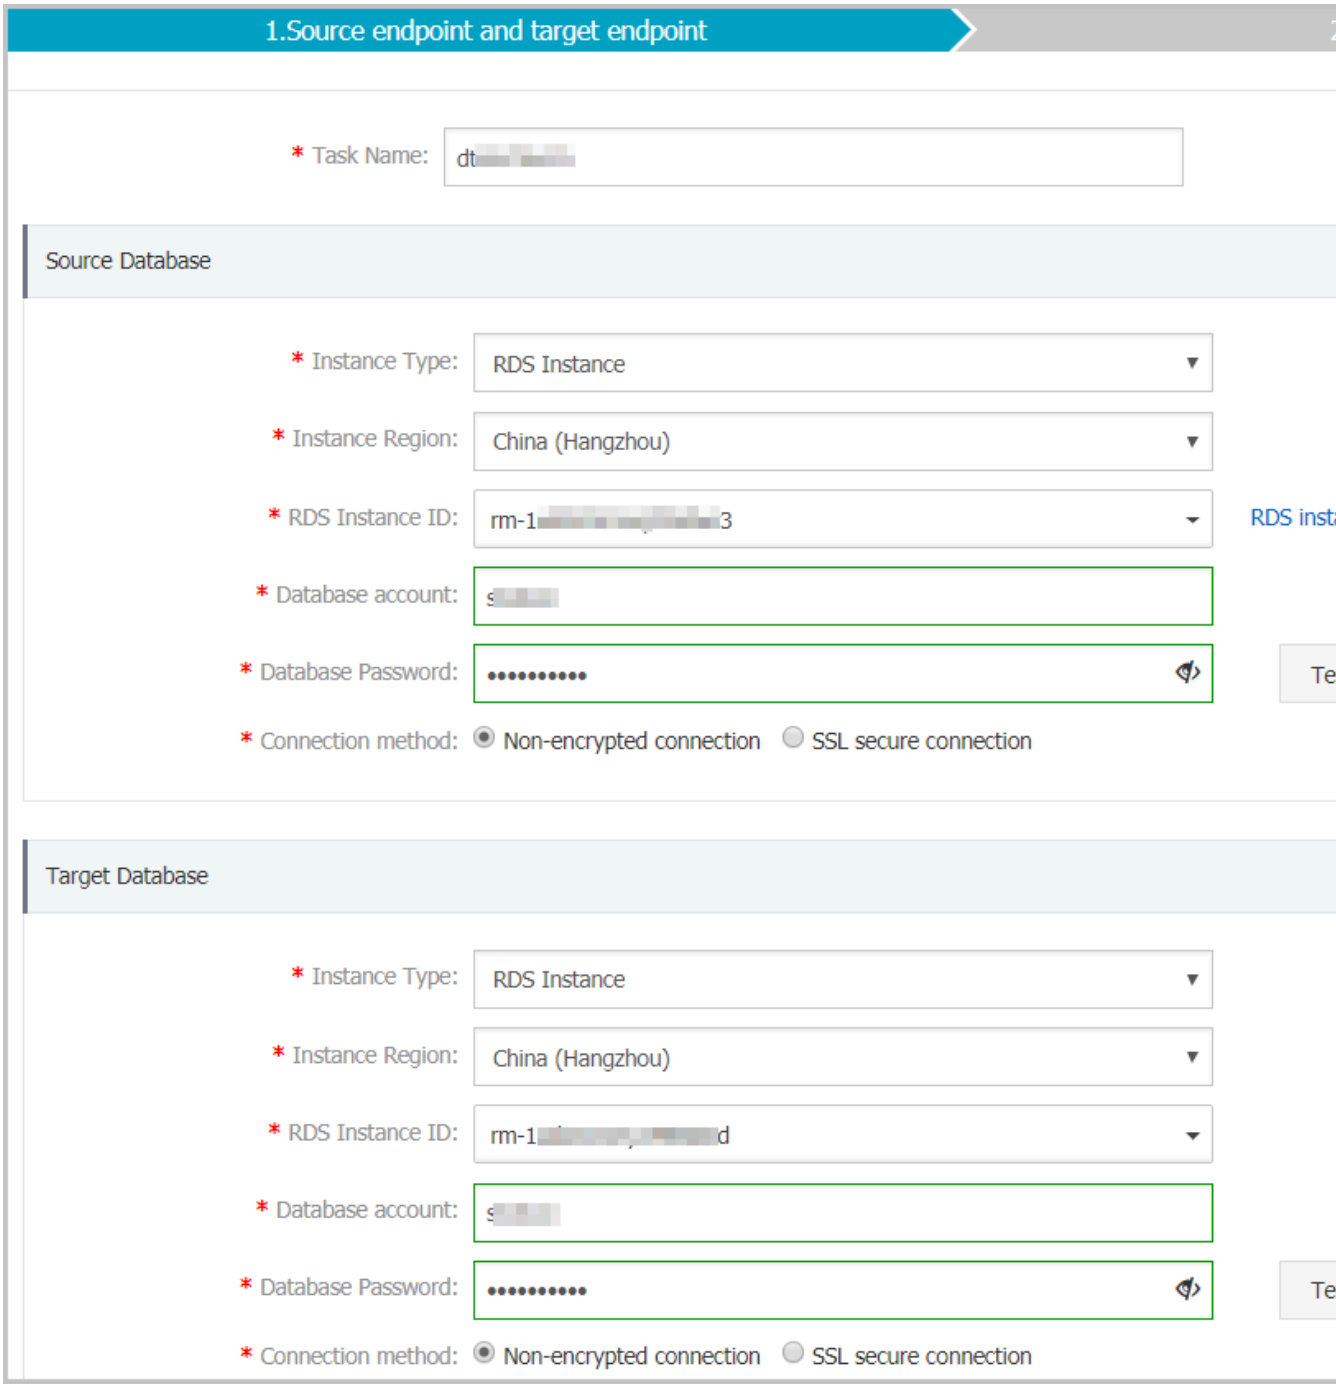

5. Click Authorized Whitelist and Enter Into Next Step.
6. Select Migrate object structure and Migrate existing data.
7. In the left pane, select objects and click > to add them to the right.

Note:

DTS will perform a data check. If an object in the target instance has the same name as an object to be migrated, the migration fails.

If an object in the target instance has the same name as an object to be migrated, do either of the following:

- In the right pane, place your mouse over an object and click Edit to modify the object name.
- Rename the object in the target instance.

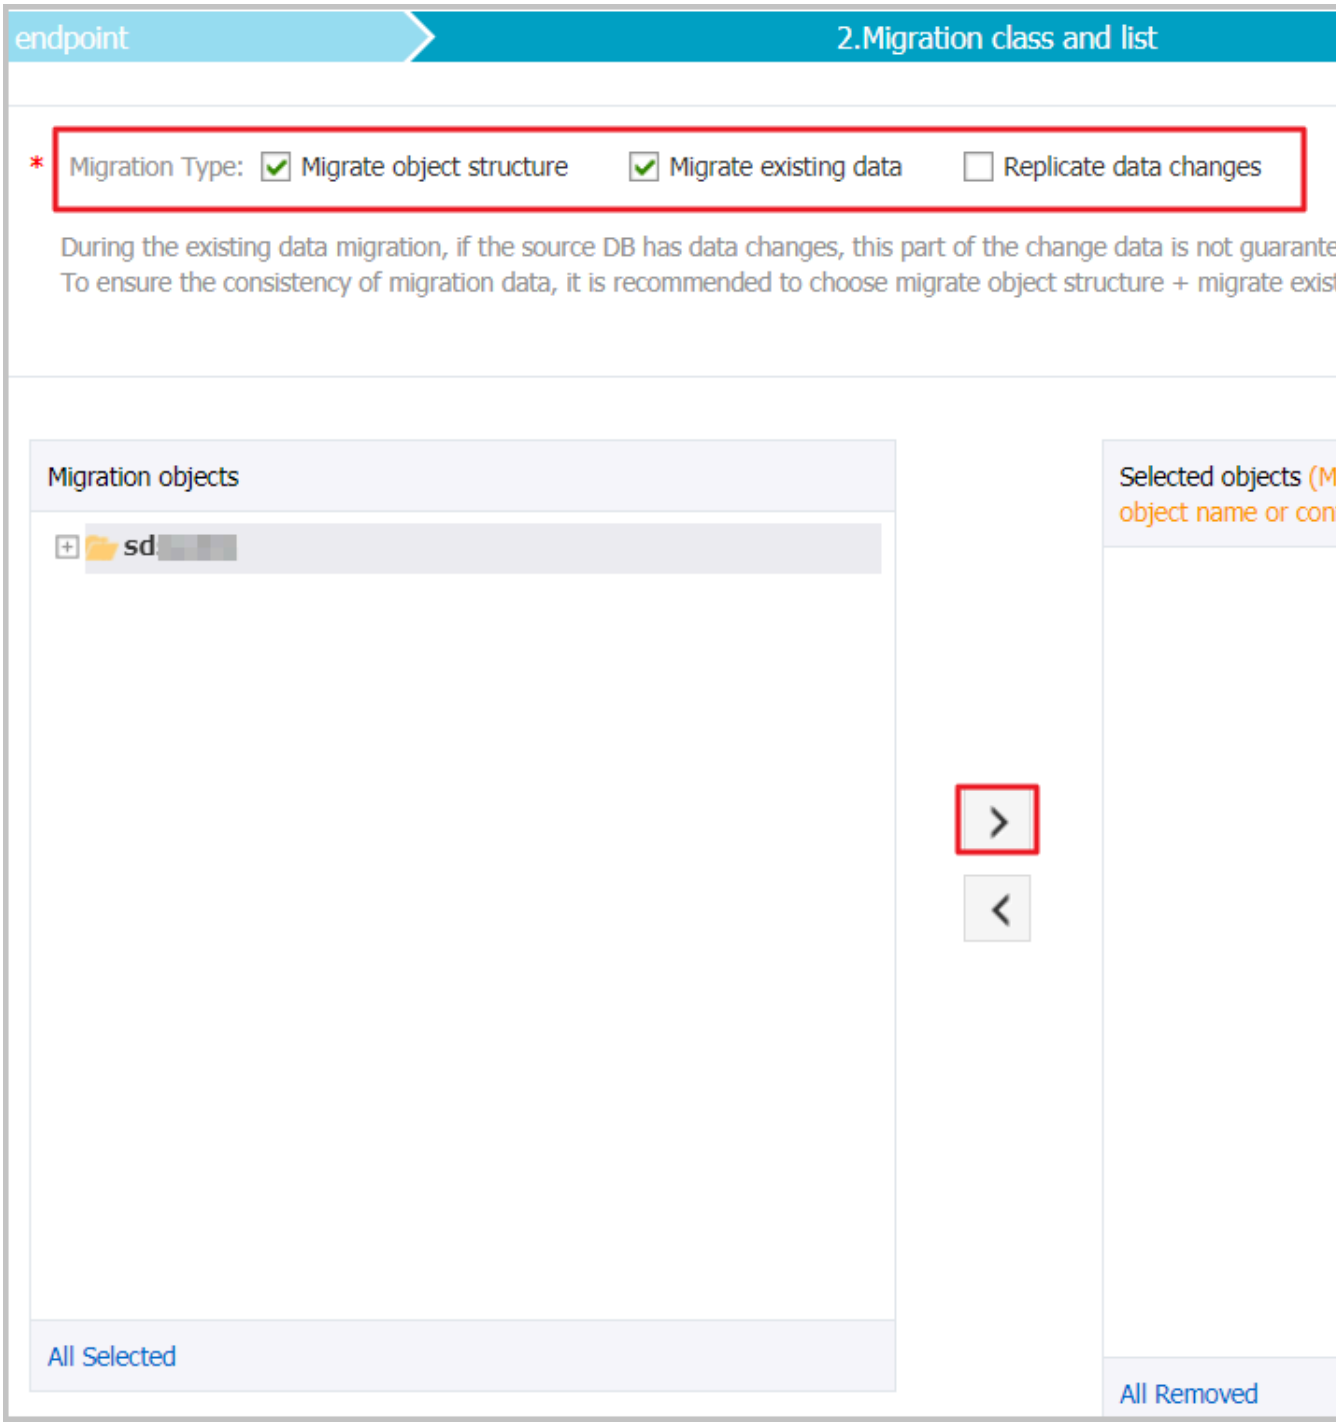

8. Click Pre-check and Start.
   - If the pre-check succeeds, go to step 11.
   - If the pre-check fails, go to step 9.
9. If the pre-check fails, click next to the failed item to view details.

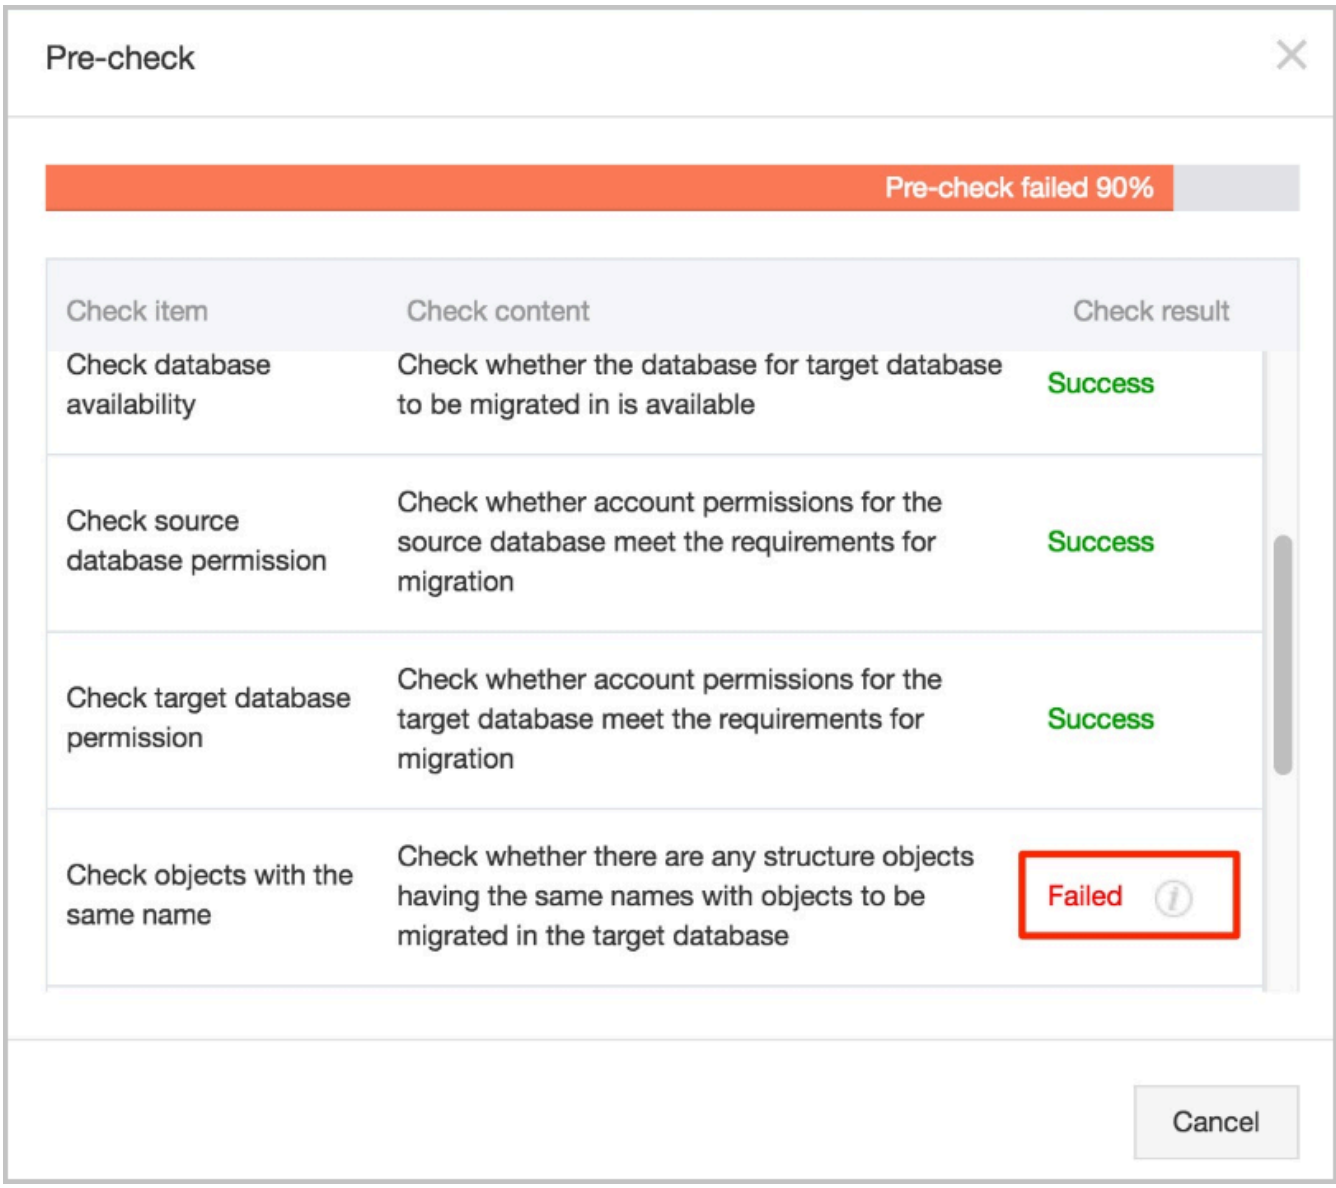

10. After fixing all problems, select the migration task in the migration task list and click Start.

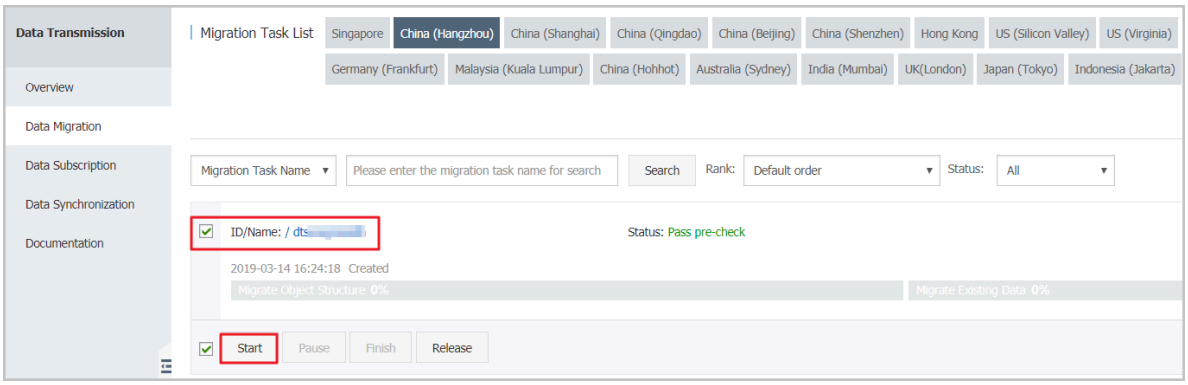

11. If the pre-check succeeds, click Next.
12. On the Confirm Purchase Configuration dialog box, confirm the configuration, select Service Terms of Data Transmission (Pay-As-You-Go), and click Buy and Start Now.

15.6 Recover data to a temporary instance (RDS for SQL Server)

> **Note:**
> This function is different from the clone instance function.

The data recovery function minimizes damage caused by incorrect operations. We recommend that you recover data to the master instance through a temporary instance. That is, recover data to a temporary instance, verify the data, and then migrate the data to the master instance. This avoids the impact of data recovery on the master instance.

Prerequisites

- The instance type is one of the following:
  - SQL Server 2012 Enterprise Basic Edition
  - SQL Server 2012/2016 Web
  - SQL Server 2008 R2
- The instance has data backups.
- To recover data to a point in time, the instance must also has log backs.

Attentions

- Creating a temporary instance does not affect the master instance.

- The temporary instance inherits the account and password of the backup file.
- The network type of the temporary instance is classic network.
- A master instance can have only one temporary instance at a time. Before creating a temporary instance, delete the existing temporary instance of the master instance.
- The temporary instance is free of charge, but will be released automatically 48 hours after being created.

Create a temporary instance

1. Log on to the *RDS console* and select the region where the target instance is located.
2. Click the ID of the target instance to go to the Basic information page.
3. Click Backup and Recovery in the left-side navigation pane.
4. Click the Temporary Instance tab.
5. Select a point in time for recovery and click Create Temporary Instance.
6. In the displayed dialog box, click OK.
7. Go back to the Instances page.

Recover data from the temporary instance to the master instance

1. After the temporary instance is created successfully, click the ID of the master instance to go to the Basic information page.
2. Click Create Data Migration Task in the upper right corner to go to the *Data Transmission Service console*.
3. Click Data migration in the left-side navigation pane.
4. Click Create migration task.

5. Set parameters.
   - Task name : A default task name is generated. You can modify it so that you can identify it more easily later.
   - Source database information:
     - Instance type : Select RDS instance.
     - Instance region : Select the region where the master instance is located.
     - RDS instance ID : Select the ID of the temporary instance.
     - Database account : It is the same as the account name of the master instance. Make sure that this account has read and write permissions on the data to be migrated.
     - Database password : It is the same as the account password of the master instance.
   - Target database information:
     - Instance type : Select RDS instance.
     - Instance region : Select the region where the master instance is located.
     - RDS instance ID : Select the master instance that has a temporary instance.
     - Database account : Enter the account name of the master instance. Make sure that this account has read and write permissions on the data to be migrated.
     - Database password : Enter the account password of the master instance.
6. Click Authorization whitelist and enter into next step.
7. Select the migration type.
8. In the left pane, select the objects to be migrated and click > to add them to the right pane. If you want to modify the name of a migrated object in the target database, you can hover the mouse over the database that needs to be modified in the Selected objects pane and click the displayed Edit button.
9. Click Pre-check and start.

10. If the pre-check fails, click ! next to the failed check item to view detailed failure information, and perform troubleshooting accordingly. After the troubleshooting, find the migration task in the Migration task list page and restart the pre-check.
11. After the pre-check is passed, click OK to start the migration task.

16 Typical applications

16.1 Cached data persistence

ApsaraDB RDS can be used together with ApsaraDB Memcache and Redis to form storage solutions with high throughput and low delay. This document describes the cached data persistence solution based on the combined use of RDS and Memcache.

Background information

Compared with RDS, Memcache and Redis have the following features:

- Quick response: The request delay of ApsaraDB Memcache and Redis is usually within several milliseconds.
- The cache area supports a higher Queries Per Second (QPS) than RDS.

System requirements

- bmemcached (with support for SASL extension) has been installed in the local environment or ECS.

  bmemcached download address: Click *Here* to download.

  The bmemcached installation command is as follows:

  ```
  pip install python-binary-memcached
  ```

- Python is used as an example. Python and pip must be installed in the local environment or ECS.

Sample code

The following sample code realizes the combined use of ApsaraDB RDS and Memcache:

```
/usr/bin/env python
import bmemcached
Memcache_client = bmemcached.Client(('ip:port'), 'user',
'passwd')
# Search for a value in ApsaraDB Memcache
res = os.client.get('test')
if res is not None:
    return res # Return the value found
else:
   # Query RDS if the value is not found
    res = mysql_client.fetchone(sql)
     Memcache_client.put('test', res) # Write cached
data to ApsaraDB for Memcache
```

```
    return   res
```

16.2 Multi-structure data storage

OSS is a cloud storage service provided by Alibaba Cloud, featuring massive capacity, security, low cost, and high reliability. RDS can work with OSS to form multiple types of data storage solutions.

For example, when the business application is a forum and RDS works with OSS, resources such as registered users' images and post content images can be stored in OSS to reduce the storage pressure of RDS.

Sample code

OSS works with the RDS.

1. Initialize OssAPI.

```
 from   oss . oss_api   import  *
 endpoint =” oss - cn - hangzhou . aliyuncs . com ”
 accessKeyI  d ,  accessKeyS  ecret =” your   id ”,” your   secret
”
 oss  =  OssAPI ( endpoint ,  accessKeyI  d ,  accessKeyS  ecret )
```

2. Create a bucket.

```
# Set   the   bucket   to   private - read - write
 res  =  oss . create_buc  ket ( bucket ," private ")
 print  "% s \ n % s " % ( res . status ,  res . read ())
```

3. Upload an object.

```
 res  =  oss . put_object  _from_file ( bucket ,  object , " test .
txt ")
 print  "% s \ n % s " % ( res . status ,  res . getheaders ())
```

4. Obtain the corresponding object.

```
 res  =  oss . get_object  _to_file ( bucket ,  object , "/
filepath / test . txt ")
 print  "% s \ n % s " % ( res . status ,  res . getheaders ())
```

In the ECS application code, RDS stores the ID of each user, and OSS stores the avatar resource of the user. The Python code is as follows:

```
/ usr / bin / env   python
 from   oss . oss_api   import  *
 endpoint =” oss - cn - hangzhou . aliyuncs . com ”
 accessKeyI  d ,  accessKeyS  ecret =” your   id ”,” your   secret ”
 oss  =  OssAPI ( endpoint ,  accessKeyI  d ,  accessKeyS  ecret )
 User_id  =  mysql_clie  nt . fetch_one  ( SQL ) #  Search   for
 user_id   in   RDS
```

```
# Obtain and download the user avatar to the
 correspond ing path
 oss . get_object _to_file ( bucket , object , your_path / user_id
 +'. png ')
# Process the uploaded user avatar
 oss . put_object _from_file ( bucket , object , your_path /
 user_id +'. png ')
```