Alibaba Cloud ApsaraDB for MySQL

User Guide

Issue: 20180807

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.	 **Note:** 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]*
{} or {a|b}	It indicates that it is a required value, and only one item can be selected.	`swich` *{stand | slave}*

Contents

1 Preface

Overview

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

This document describes RDS features and functions and further explains the procedure to configure RDS through the *RDS console* . With the help of this information you can also manage the 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 the *official website of ApsaraDB for RDS*.

Declaration

Some product 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 as per the product upgrade and many other respective factors. You must first verify the document with your latest corresponding software version.

Consideration

RDS includes multiple types of databases. This document takes the MySQL database as an example to describe the features and usage of all the RDS products. Some types of databases may not include 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 specificat

ion 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 datacenter. A zone is a physical area that has independent power supply and networks. For more information, see *Alibaba Cloud Global Infrastructure*.

General terms

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 and whose database type is MySQL.

2 Quick start

If you use RDS for the first time, see the following *Cite LeftQuick StartCite Right* 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 *Cite LeftQuick StartCite Right*, see the *Cite LeftUser GuideCite Right*.

Database engines

ApsaraDB for MySQL

MySQL is the world's most popular open source database. As an important part of LAMP, 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 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 a visual interface.

Powered by high-availability architecture and anytime data recovery capabilities, ApsaraDB for SQL Server provides strong support for a variety of enterprise applications. It also covers Microsoft's licensing fee without any additional cost required.

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, what sets PostgreSQL apart is that its 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 verification. It also integrates a number of important functions, including high availability and backup recovery, to help mitigate your operation and maintenance 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 controls, and it continues to be updated and improved regularly.

RDS for PPAS currently supports version 9.3.

3 Introduction to RDS console

3.1 Logon and logout

Log on

1. Log on.
   - If you are using an Alibaba account, click *here* to log on.
   - If you are using a sub-account, click *here* to log on.
2. Access the *RDS console*.
3. Select a region.

   The list of instances in the region is displayed.

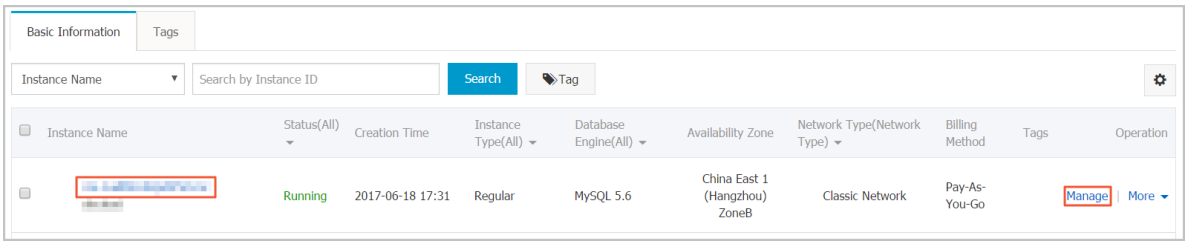

Log out

To log out from the RDS console, place the cursor over your account information at the upper right corner of the console, and click **Sign out**.

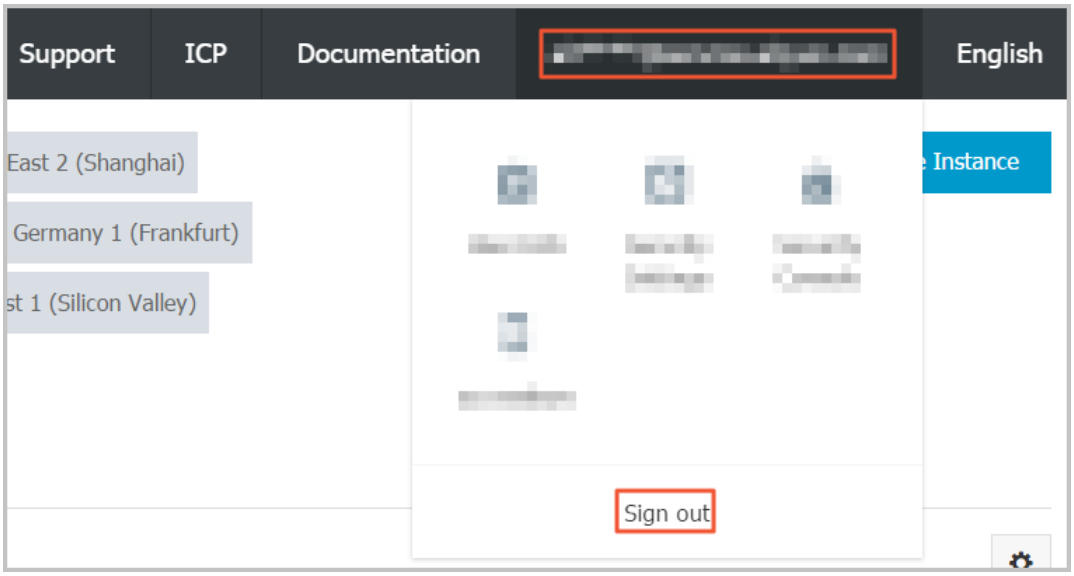

Introduction to the RDS console

- *Introduction to the RDS console (MySQL)*
- *Introduction to the RDS console (SQL Server)*
- *Introduction to the RDS console (PostgreSQL)*
- *Introduction to the RDS console (PPAS)*

3.2 The homepage of RDS console

The *RDS console* is a web application used to manage the RDS instances. The console has intuitive user interface, through which you can create instances, create databases, create accounts, set network, set connection, and so on. For the different instance types, the information displayed on the console is different.

The RDS console is part of the Alibaba Cloud console. For information about common settings and basic operations on the Alibaba Cloud console, see *Alibaba Cloud console*. This article describes the home page of the RDS console. The actual interface may vary slightly.

Log on to the *RDS console* to enter the **Instance List** page, which is the homepage of the RDS console, as shown in the following figure.

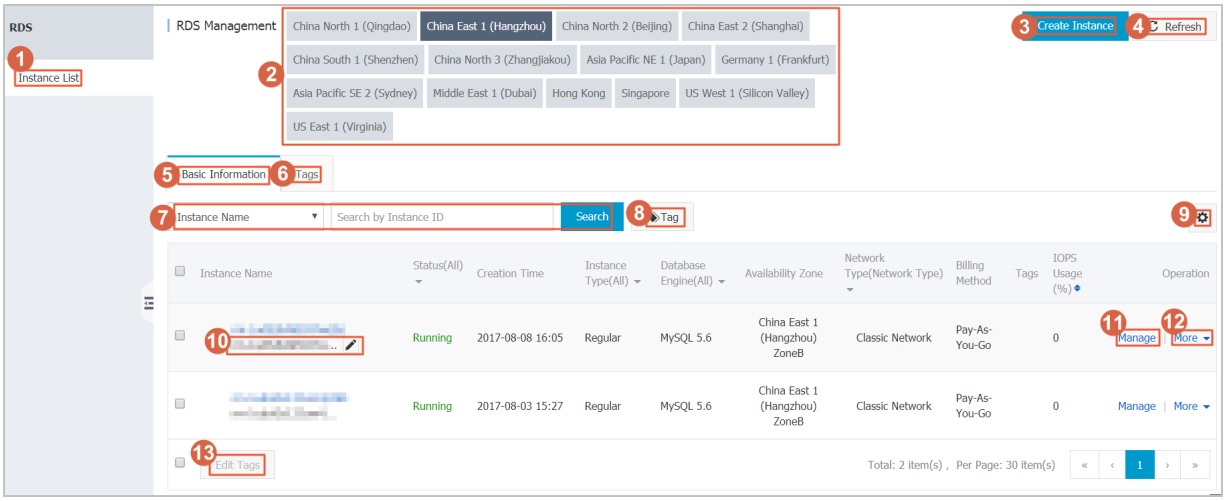

Parameter description:

- 1: The **Instances** page is also the homepage of the RDS console and displays all the instances under the same account.
- 2: Region name, select a certain region and all the instances in this region are displayed in the instance list.
- 3: Instance creation portal.
- 4: Refresh the instance list information.

- 5: List of all instances in a region.
- 6: The page shows a list of tags added by users.
- 7: Instance search field.
- 8: If you add a tag in an instance, the tag content is displayed here. For more information on how to add a tag, see *Create tags*.
- 9: Set the columns to be displayed in the instance information list.
- 10: The remark name of the instance, and it is the same with the instance ID by default. Click the edit icon to modify the remark name if needed.
- 11: Click the button to enter the management details page of the instance to view basic information, set the network and connection mode, and create a database.
- 12: Shortcut keys for some operations, such as editing tags. Click **More** to show more operations. Different operations are available for instances of different types. See the actual interface when using this document.
- 13: Edit tags in batches.

3.3 The instance management interface for MySQL

This document introduces the query information and executable operations supported by the RDS console for a MySQL instance.

Log on to the instance management interface

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 **Manage** in the corresponding **Operation** column to enter the instance management interface.

Overview of instance management interface

The following table lists the query information and executable operations supported by the RDS console for MySQL instance. As MySQL instances of different versions support different operations, information displayed on the console may vary accordingly. See the actual interface when using this document.

Navigation pane	Block name	Description	Links of common operations
Operation buttons at the top of the page		You can migrate the database, and restart	• *Restart instance* • *Back up RDS data*

		or back up an instance .	
Basic Information	Basic Information	You can view the basic information of an instance. For example, the instance ID, region, and zone, instance type, intranet and Internet addresses, intranet and Internet ports, and perform migration across zones.	*Instance migration across zones*
	Instance Distribution	You can query the number of read-only and temporary instances under the primary instance, and the operation of adding read-only instances, adding temporary instances and so on.	*Create a read-only instance*
	Operating Status	You can view the running status, billing method, and creation time of an instance , and release an instance, renew a subscription instance and so on.	• *Release an instance* • *Manually renew a subscription instance*
	Configuration Information	You can view the instance types, CPU , database type and version, database memory, and the maximum number of connections, and upgrade the database version, set the maintenance period and so on.	• *Upgrade database* • *Configure the maintenance period*

	Resource Information	You can view the storage space and backup usage of an instance.	
Account Management	Account List	You can do the following: View all accounts under an instance; Create accounts, master account, or initial account; Change the account password, delete an account, and modify the account permissions.	• *Create account* • *Create account and database for MySQL 5.7 High-availability Edition/5.5/5.6 instances* • *Create account and database for MySQL 5.7 Basic Edition* • *Create master account* • *Reset instance password* • *Change account permissions*
	Service Account Privileges	When an Alibaba Cloud engineer provides technical support, you must authorize the engineer's service account to view or modify the instance configurations and view the table structure, index, and SQL statements.	*Authorize a service account*
Database Management		You can view the databases information under an instance, and create and delete databases.	• *Create account and database for MySQL 5.7 High-availability Edition/5.5/5.6 instances* • *Create account and database for MySQL 5.7 Basic Edition* • *Delete database*

Database Connection	Instance Connection	You can view the network type, access mode, intranet address, and port of an instance, change the network type, modify the connection address, and to apply for and release an intranet/Internet address.	• *Set access mode* • *Set network type* • *Set intranet and the Internet addresses*
Monitoring and Alarms	Monitoring	You can view the monitoring information, such as the CPU and memory usage, disk space usage, and IOPS, and set the monitoring frequency.	*Set monitoring frequency*
	Alarms	You can set the alarm rules and view the status of monitoring items and the alarm contact.	*Set monitoring rules*
Security Controls	Whitelist settings	You can view the whitelist information of an instance, modify the whitelist, and add a whitelist group.	*Set whitelist*
	SSL	You can view an SSL certificate, set SSL, and download a certificate.	*Set SSL encryption*
	TDE	You can view the status of Transparent Data Encryption (TDE) and activate the TDE.	*Set Transparent Data Encryption*
Instance Availability	Availability Information	You can view the instance zone type, instance availability, data replication mode	• *Switch master/ slave instance*

		, and ID of the master/ slave node, switch the master/slave instances , and modify the data replication mode.	• *Modify the data replication mode*
	Zone Architecture Diagram	You can view the structural diagrams of a single zone and multi-zone instance.	
Log Management	Error Log	You can view SQL statements that are incorrectly executed in the database within a month.	*Log management*
	Slow SQL Log Details	You can view SQL statements whose execution period exceeds one second in the database within a month, and deduplicate similar statements.	
	Slow SQL Log Summary	Collects SQL statements whose execution period exceeds one second in the database within a month, provides the analysis report for slow query logs , and allows you to download the statistics list.	
Backup and Recovery	Backup List	You can view the data backup list, recover data to the master instance, and delete and download the backup data.	• *Recover data to the master instance through a temporary instance* • *Download RDS data and log backup*
	Binlog List	You can view and download binlog files.	

	Temporary Instance	You can create temporary instances which can be used to restore the lost data.	
	Backup Settings	You can modify a backup policy and view the backup policies, such as the data backup retention time, backup cycle, and backup time.	*Back up RDS data*
Parameter Settings	Modifiable Parameters	You can view the parameter values of an instance, modify the parameter values, and import and export the parameters.	*Set parameters through RDS console*
	Modification History	You can view parameter modification records.	

3.4 The instance management interface for SQL Server

This document introduces the query information and executable operations supported by the RDS console for an SQL Server instance.

Log on to the instance management interface

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 **Manage** in the corresponding **Operation** column to enter the instance management interface.

Overview of instance management interface

The following table lists the available information and actions that can be performed. The actual interface may vary slightly with the SQL Server version.

Navigation pane	Block name	Description	Links of common operations
Operating area at the top		Allows you to migrate the database, restart	• *Restart instance*

Navigation pane	Block name	Description	Links of common operations
		, and back up an instance.	• *Back up RDS data*
Basic Information	Basic Information	Allows you to view the basic information of an instance. For example, the instance ID, region, and zone, instance type, intranet, and Internet addresses, intranet, and Internet ports, and perform migration across zones.	*Instance migration across zones*
	Instance Distribution	Allows you to check the number of temporary instances under the master instance and add a temporary instance.	
	Operating Status	Allows you to view the running status, billing method, and creation time of an instance, release an instance, and renew a subscription instance.	• *Release an instance* • *Manually renew the subscription instance*
	Configuration Information	Allows you to view the instance types, CPU, database type and version, database memory, and the maximum number of connections, and set the maintenance period.	• *Configure the maintenance period*
	Resource Information	Allows you to view the storage space and	

Navigation pane	Block name	Description	Links of common operations
		backup usage of an instance.	
Account Management	Account List	Allows you to do the following: View all accounts under an instance; Create accounts, master account, or initial account; Change the account password, delete an account, and modify the account permissions.	• *Create database and account for SQL Server 2008 R2* • *Create database and account for SQL Server 2012 and 2016* • *Reset instance password* • *Change account permissions*
	Service Account Privileges	When an Alibaba Cloud engineer provides technical support, you must authorize the engineer's service account to view or modify instance configurations and view the table structure, index, and SQL statements.	*Authorize a service account*
Database Management		Allows you to view the databases information under an instance, create, and delete a database.	• *Create database and account for SQL Server 2008 R2* • *Create database and account for SQL Server 2012 and 2016* • *Delete database*
Database Connection	Instance Connection	Allows you to view the network type, access mode, intranet and Internet addresses,	• *Set access mode* • *Set network type*

Navigation pane	Block name	Description	Links of common operations
		intranet and Internet ports, and server name of an instance, change the network type, modify the connection address, and to apply for and release an intranet/ Internet address.	• *Set intranet and the Internet addresses*
Monitoring and Alarms	Monitoring	Allows you to view the monitoring information, such as the CPU and memory usage, disk space usage, and IOPS, and set the monitoring frequency.	*Set monitoring frequency*
	Alarms	Allows you to view the status of monitoring items and cloud account alert contact, and set the alarm rules.	*Set alarm rules*
Security Controls	Whitelist Settings	Allows you to view the whitelist information of an instance, modify the whitelist, and add a whitelist group.	*Set whitelist*
	SSL	Allows you to view an SSL certificate, set SSL, and download a certificate.	*Set SSL encryption*
	TDE	Allows you to view the status of Transparent Data Encryption (TDE) and activate the TDE.	*Set Transparent Data Encryption*
Instance Availability	Availability Information	Allows you to view the instance zone type	*Switch master/slave instance*

Navigation pane	Block name	Description	Links of common operations
		, instance availabili ty, data replication mode, and number of the master/slave database, and switch the master/slave instances.	
	Zone Architecture Diagram	Allows you to view the structural diagrams of the single zone and multi-zone instance.	
Log Management	Error Log	Allows you to view SQL statements that are incorrectl y executed in the database within a month.	*Log management*
	Slow SQL Log Summary	Collects SQL statements whose execution period exceeds one second in the database within a month, provides the analysis report for slow query logs , and allows you to download the statistics list.	
Backup and Recovery	Backup List	Allows you to view the data backup list, and recover data to the master instance.	
	Temporary Instance	Allows you to create temporary instances.	
	Backup Settings	Allows you to view the backup policies, such as the data backup retention time, backup	*Back up RDS data*

Navigation pane	Block name	Description	Links of common operations
		cycle, and backup time , and modify a backup policy.	
Parameter Settings	Modifiable Parameters	Allows you to view the parameter values of an instance, modify the parameter values, and import and export the parameters.	*Set parameters through RDS console*
	Modification History	Allows you to view parameter modification records.	

3.5 The instance management interface for PostgreSQL

This document introduces the query information and executable operations supported by the RDS console for the PostgreSQL instance.

Log on to the instance management interface

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 **Manage** in the corresponding **Operation** column to enter the instance management interface.

Overview of the instance management interface

The following table lists the available information and actions that can be performed.

Navigation pane	Area name	Description	Links of common operations
Operation buttons at the top of the page		This area allows you to migrate the database, restart, and back up an instance.	• *Restart an instance* • *Back up RDS data*
Basic Information	Basic Information	This area allows you to view the the basic information of an instance. For example	

Navigation pane	Area name	Description	Links of common operations
		, the instance ID, region, and zone, instance type, intranet and the Internet addresses, intranet and the Internet port number.	
	Instance Distribution	This area allows you to check the number of temporary instances under a master instance, and add temporary instances.	
	Operating Status	This area allows you to view the running status, billing method, and creation time of an instance, release an instance, and renew a subscription instance.	• *Release an instance* • *Manually renew the subscription instance*
	Configuration Information	This area allows you to view the instance types, CPU , database type and version, database memory, and the maximum number of connections, and set the maintenance period.	*Configure the maintenance period*
	Resource Information	This area allows you to view the storage space and backup usage of an instance.	
Account Management	Account List	This area allows you to view the account information of the instance, create an initial account, and	• *Create database and account* • *Reset instance password*

Navigation pane	Area name	Description	Links of common operations
		modify the account password.	
Database connection	Instance connection	This area allows you to view the network type, access mode, intranet and the Internet addresses, and port of an instance, change the network type, modify the connection address, and to apply for and release an intranet/ Internet address.	• *Set access mode* • *Set network types* • *Set intranet and Internet addresses*
Monitoring and Alarms	Monitoring	This area allows you to view the monitoring information, such as the CPU and memory usage, disk space usage, and IOPS, and set the monitoring frequency.	*Set monitoring frequency*
	Alarms	This area allows you to view the status of monitoring items, cloud account, and the alarm contact, and set alarm rules.	*Set monitoring rules*
Security Controls	Whitelist settings	This area allows you to view the whitelist information of an instance, modify the whitelist, and add a whitelist group.	*Set whitelist*
Instance Availability	Availability Information	This area allows you to view the instance zone type, instance availability, data replication mode, and	*Switch master/slave instance*

Navigation pane	Area name	Description	Links of common operations
		number of the master /slave database, and switch the master/ slave instances.	
	Zone Architecture Diagram	This area allows you to view the structural diagrams of the single zone and multi-zone instance.	
Log management	Error Log	This area allows you to view SQL statements that are incorrectly run in the database within a month.	*Log management*
	Slow SQL Log Details	This area allows you to view SQL statements whose running period exceeds one second in the database within a month, and deduplicate similar statements.	
Backup and Recovery	Backup List	This area allows you to view the data backup list, and download the backup data.	*Download RDS data and log backup*
	Temporary instance	This area allows you to create temporary instances.	
	Archive List	This area allows you to view the detailed list of archived logs, and download the archived logs.	

Navigation pane	Area name	Description	Links of common operations
	Backup Settings	This area allows you to modify a backup policy and view the backup policies, such as the data backup retention time, backup cycle, and backup time.	*Back up RDS data*

3.6 The instance management interface for PPAS

This document introduces the query information and executable operations supported by the RDS console for the PPAS instance.

Log on to the instance management interface

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 **Manage** in the corresponding **Operation** column to enter the instance management interface.

Overview of instance management interface

The following table lists the query information and executable operations supported by the RDS console for PPAS instance. As PPAS instances of different versions support different operations, information displayed on the console may vary accordingly. See the actual interface when using this document.

Navigation pane	Block name	Description	Links of common operations
Operation buttons at the top of page		Allows you to migrate the database, restart, and back up an instance.	• *Restart an instance* • *Back up RDS data*
Basic Information	Basic Information	Allows you to view basic information of an instance. For example, the instance ID, region, and	

Navigation pane	Block name	Description	Links of common operations
		zone, instance type, intranet, and Internet addresses, intranet , and Internet port number, and so on.	
	Instance Distribution	Allows you to check the number of temporary instances under a master instance, and add temporary instances.	
	Operating Status	Allows you to view the running status, billing method, and creation time of an instance, release an instance, and renew a subscription instance.	• *Release an instance* • *Manually renew a subscription instance*
	Configuration Information	Allows you to view the instance types, CPU, database type and version, database memory, and the maximum number of connections, and set the maintenance period.	*Configure the maintenance period*
	Resource Information	Allows you to view the storage space and backup usage of an instance.	
Account Management	Account List	Allows you to create the initial account, view the initial account information, and modify the account password.	*Reset instance password*

Navigation pane	Block name	Description	Links of common operations
Database Connection	Instance Connection	Allows you to view the network type, access mode, intranet, and Internet addresses, and port of an instance, change the network type, modify the connection address, and to apply for and release an intranet/ Internet address.	• *Set access mode* • *Set network type* • *Set intranet and Internet addresses*
Monitoring and Alarms	Monitoring	Allows you to view the monitoring information, such as the CPU and memory usage, disk space usage, and IOPS, and set the monitoring frequency.	*Set monitoring frequency*
	Alarms	Allows you to view the status of monitoring items and cloud account alert contact, and set the alarm rules.	*Set alarm rules*
Security Controls	Whitelist settings	Allows you to view the whitelist information of an instance, modify the whitelist, and add a whitelist group.	*Set whitelist*
Instance Availability	Availability Information	Allows you to view the instance zone type, instance availability, data replication mode, and number of the master/slave database, and switch the master/slave instances.	*Switch master/slave instance*

Navigation pane	Block name	Description	Links of common operations
	Zone Architecture Diagram	Allows you to view the structural diagrams of the single zone and multi-zone instance.	
Log Management	Error Log	Allows you to view SQL statements that are incorrectly executed in the database within a month.	*Log management*
	Slow SQL Log Details	Allows you to view SQL statements whose execution period exceeds one second in the database within a month, and deduplicate similar statements.	
Backup and Recovery	Backup List	Allows you to view the data backup list, and recover data to the master instance.	
	Temporary Instance	Allows you to create temporary instances.	
	Archive List	Allows you to view the details list of archived logs, and download archived logs.	
	Backup Settings	Allows you to view the backup policies, such as the data backup retention time, backup cycle, and backup time, and modify a backup policy.	*Back up RDS data*

4 Instance management

4.1 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* and select the target instance.
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**.

4.2 Change billing method

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

Attention

- Think twice before such 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 and cannot be 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 on the *Orders* page and you cannot purchase new instances or change the 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 not enough. 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.

Prerequistes

- 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 the order only when the instance status restores to Running.

- No unfilled instance order for billing method resetting (i.e. new subscription order).

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**.

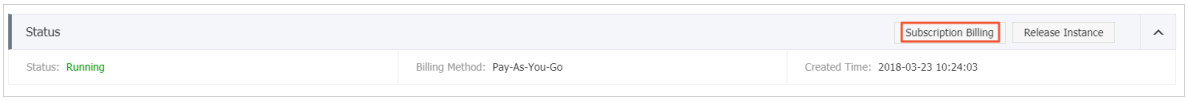

5. Select the subscription length.
6. Click **Pay Now** and pay for the order.

4.3 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 the potential interruptions for production during maintenance. RDS performs regular maintenance within the period you have configured.

Background information

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

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

To guarantee stability during 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 the database is not affected. However, apart from account management, database management, adding IP addresses to the whitelist, and other services associated with changes (such as common operations including upgrade, degrade, and restart) are unavailable on the console of this instance. Query services such as performance monitoring are available.

When the maintenance period preset by an instance begins, transient disconnection occurs once or twice on the instance during this period. You must make sure that the application program supports the reconnection policy. After transient disconnection, the instance restores to the normal state.

Procedure

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

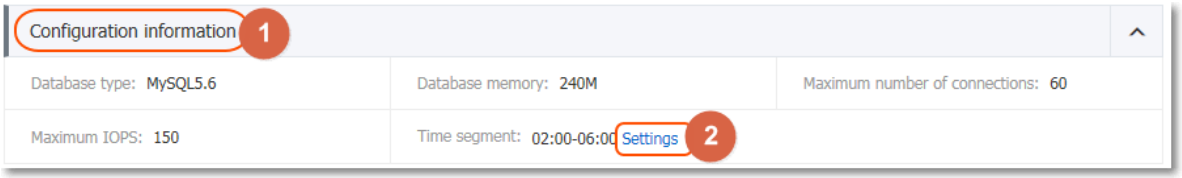

4. Select the maintenance period and click **Save**, as shown in the following figure.

> **Note:**
> Note: Time segment is the time in Beijing.

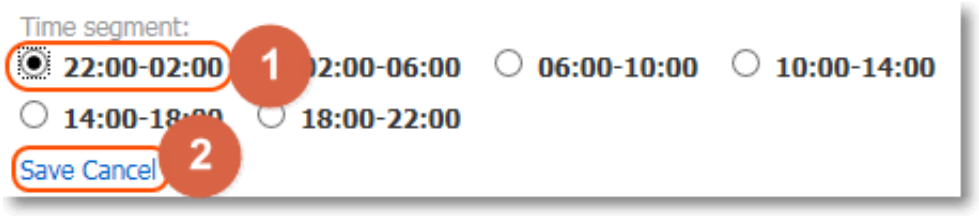

4.4 Instance migration across zones

If the zone in which the instance is located is in full load or the instance performance may be 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 be off-peak hours. This article describes the 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 RDS instances. A multi-zone is a physical zone created by combining 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 interface). No extra charge is collected for the use of the multi-zone.

If the zone in which the instance is located is in full load or the instance performance may be 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 have the following options for migration 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 backup database, the two databases are randomly allocated in the multi-zone. For example, when an instance having a master database and a backup database is migrated from Zone A to Multi-zone 1 (Zone B + Zone C), if the master database is allocated to Zone B, the backup database is allocated to Zone C.
- Migrate the instance from a multi-zone to a single-zone. In this case, the master and backup databases of the instance are migrated to the same zone, and the level of disasters that the instance can withstand is lower.

Note:
Because certain network delay occurs 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 solution. 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:

Function	Can be performed
Modify the whitelist	Yes
Enable SQL audit	Yes
Set the maintainenance period	Yes
Add read-only instances	No
Add disaster-recovery instances	No
Release an instance	No
Switch to the Subscription mode	No
Change configurations	No
Create a common or master account	No

Function	Can be performed
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 address	No
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. Please 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** module, as shown in the following figure.

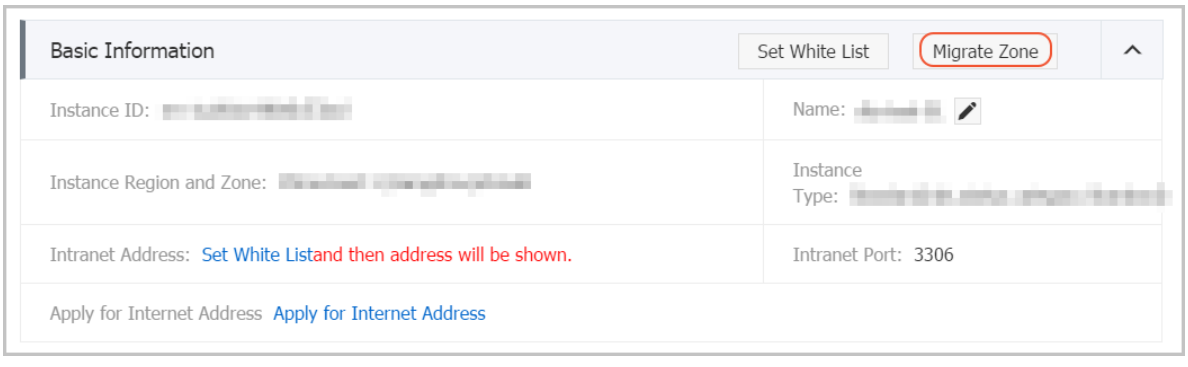

5. Select a target zone on the **Migrate Zone** page

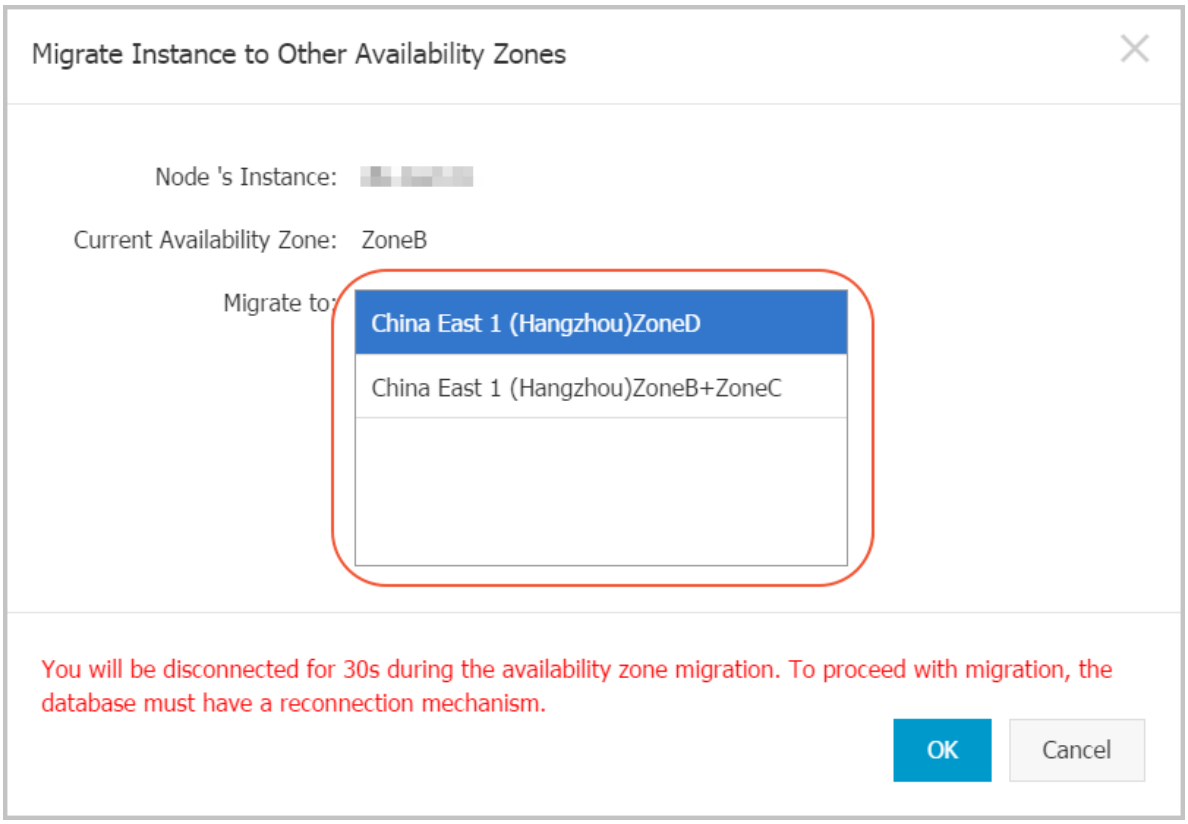

Parameter description:

- 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: You can also leave the maintenance time unchanged.

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

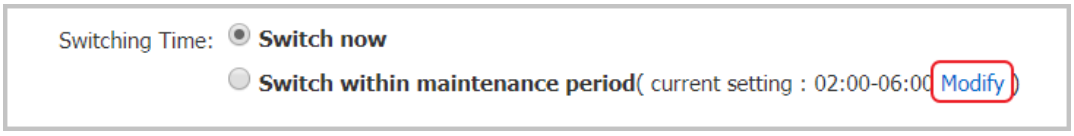

   b. At the lower left corner, select a maintenance period in the **`Configuration Information`** area and click **Save**.

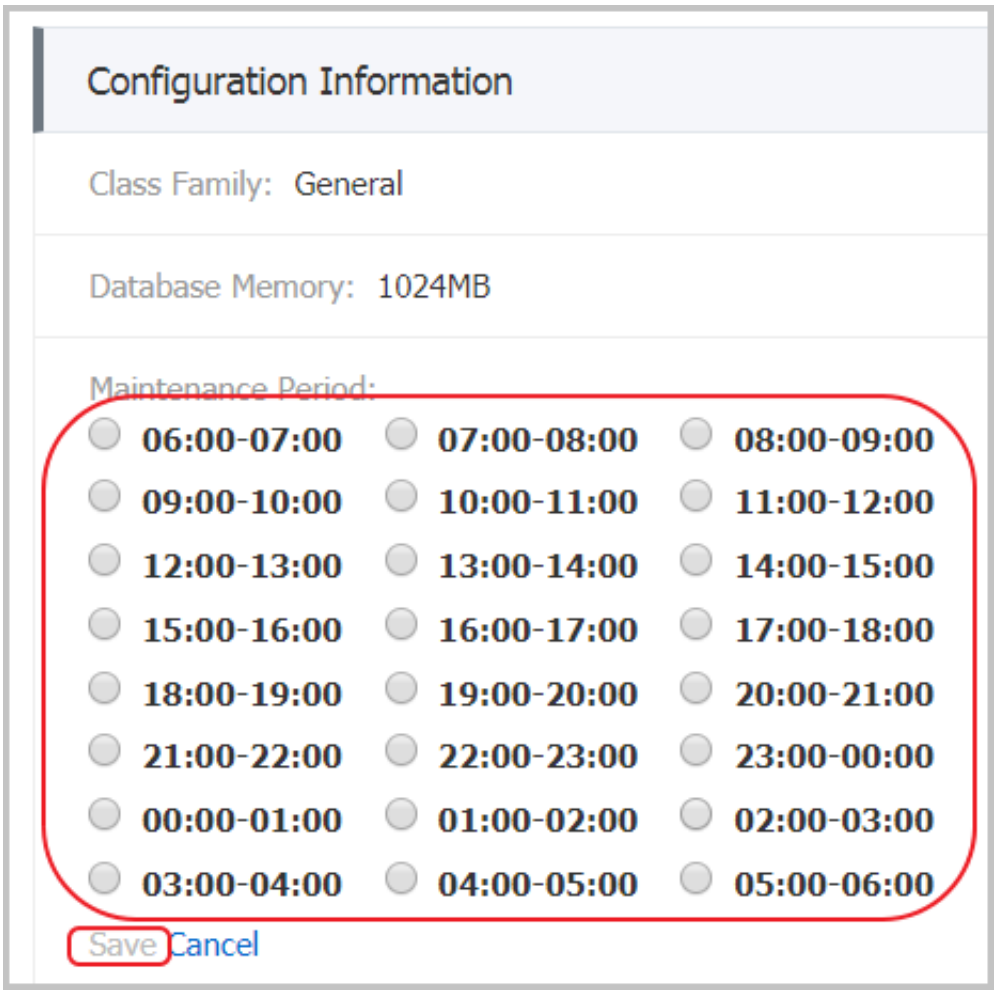

c. Go back to page for migrating the instance to another zone.

7. On the **Migrate Instance to Other Availability Zones** page, click **OK**.

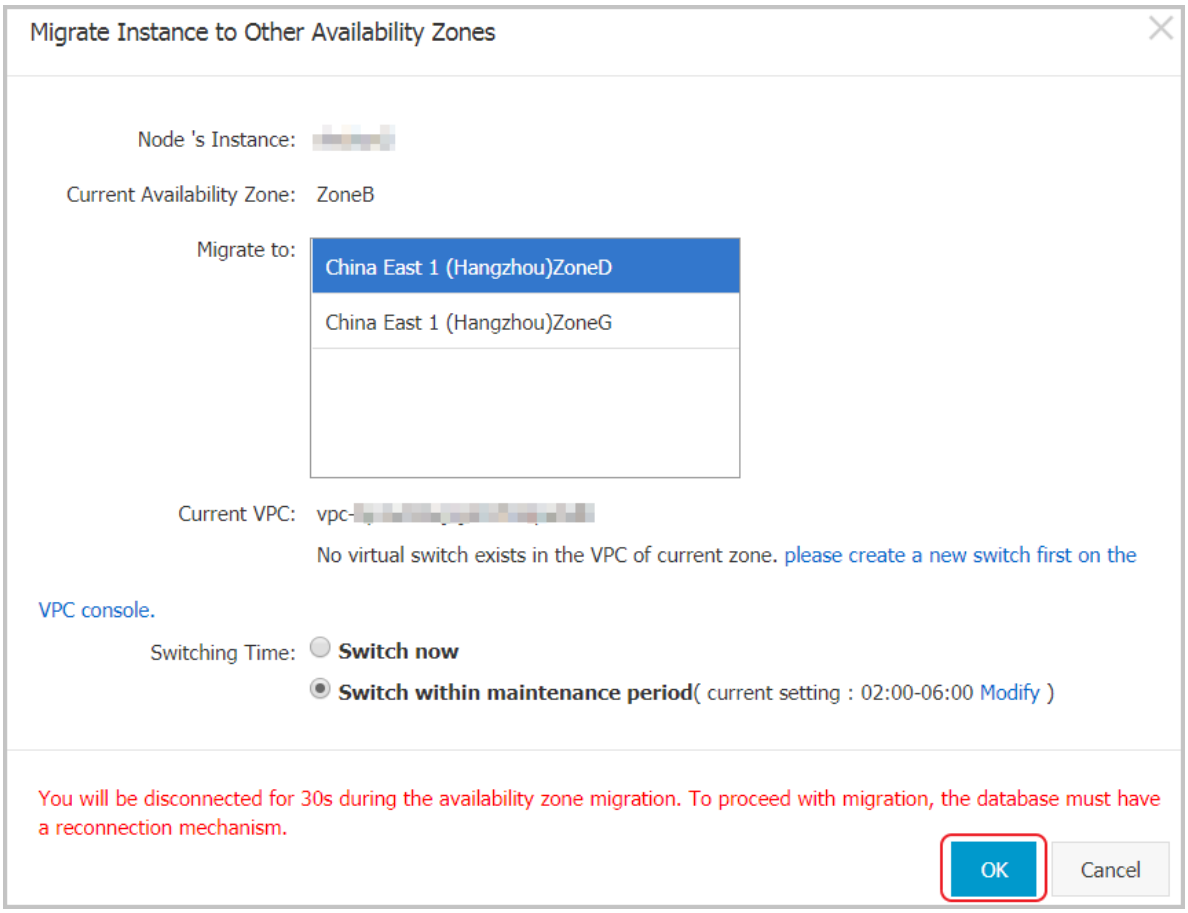

4.5 Switch master/slave instance

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

The data in the master instance synchronizes 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 handover, the original master instance becomes a backup instance, and the rack-level disaster tolerance can be realized.

This document describes how to switch the master/slave instance.

Attention

- 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 the master/slave instance may result in transient disconnection. 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 **`Availability Information`** area, click **Switch Master/Slave Instance**.
5. Select **Switch now** or **Switch within maintenance period**.

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

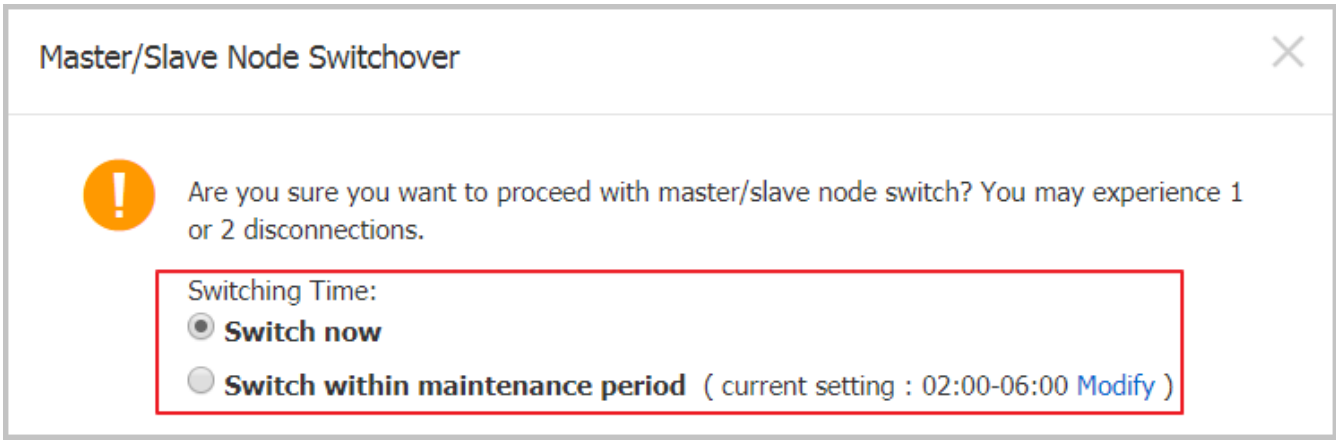

6. If necessary, you can change the maintenance period as follows:
   a. Click **Modify** to open the **Basic Information** page.

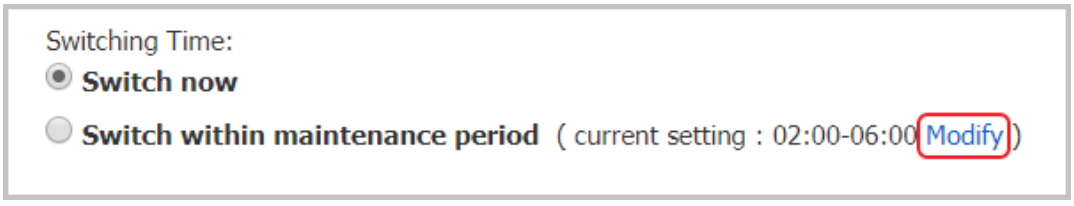

   b. In the **Configuration Information** area at the lower left corner, select a maintenance period and click **Save**.

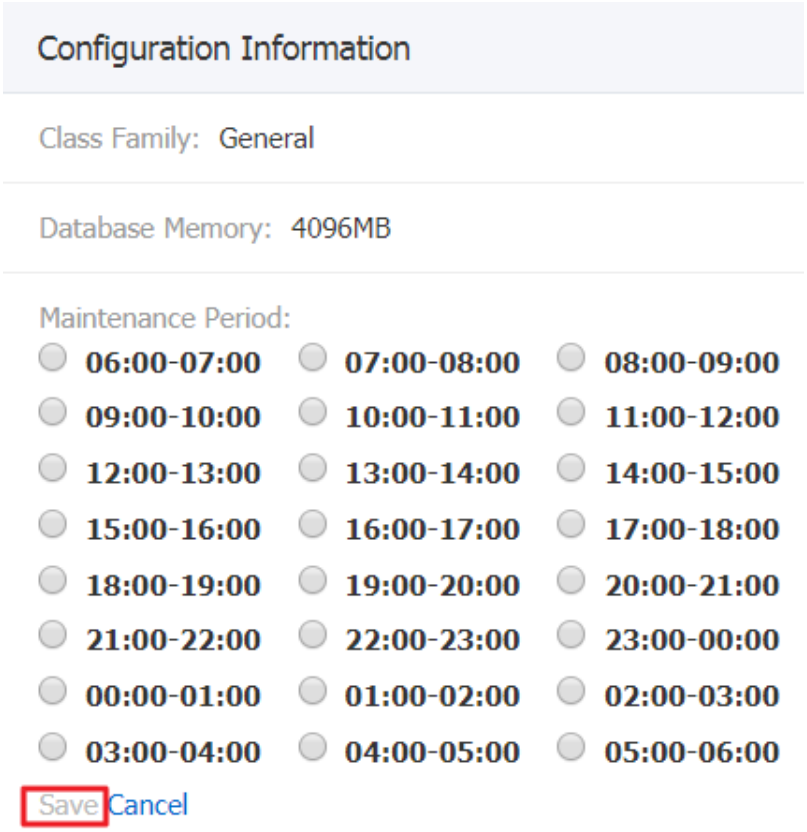

c. Go back to the page for master/slave switchover and refresh the page.

7. Click **OK**.

4.6 Modify the data replication mode

For MySQL 5.5/5.6 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.

Note:
The Finance Edition instance has one master node and multiple slave nodes. This kind of instance only supports the strong synchronous replication mode by default, which cannot be modified.

Background information

MySQL 5.5/5.6 instance supports three replication modes: sync, semi-sync and async. You can select the proper replication mode as your business needs. The differences and features of the two replication modes are shown as follows.

- Sync mode:

- When the updates initiated by application are all executed in the master node, the log is synchronously transferred to all the slave nodes. The transaction completes the commit only when most nodes (including the master node) in the cluster receive and store the log.
- Only the instance that has three or more nodes supports the strong synchronous replicatio n. In the strong synchronous replication, no matter what happens, the replication mode cannot degrade into the asynchronous replication mode.

- 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 the application until the replication mode times out and degrades to 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 length of time to restore to the sync mode depends on the implementation mode of semi-sync mode. ApsaraDB for MySQL 5.5 is different from ApsaraDB for MySQL 5.6 in this regard.
- Async mode: The application initiates an update (including the Add, Delete, and Modify operations) request. After completing the corresponding operation, the master node immediatel y 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 primary 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.

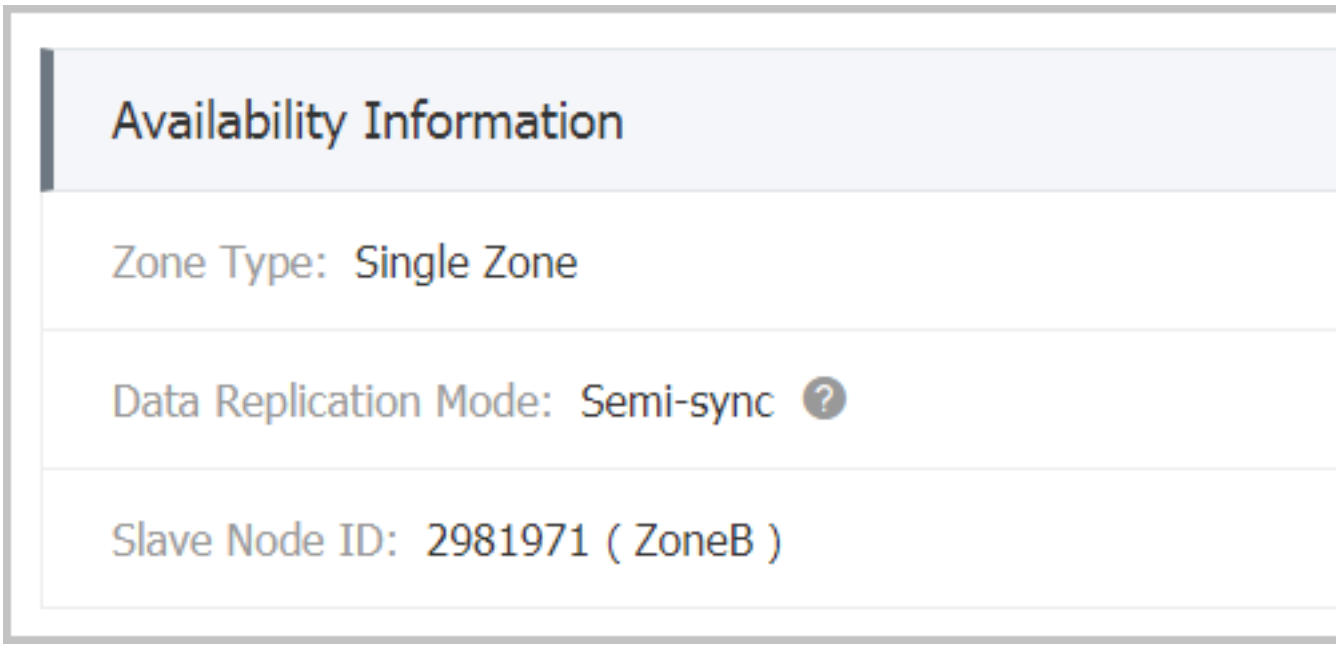

6. On the **Modify Data Replication Mode** page, select a data replication mode, as shown in the following figure.

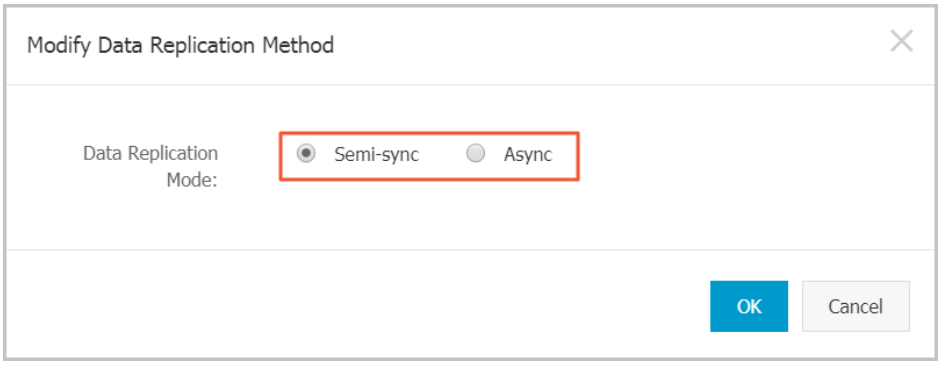

7. Click **OK**.

4.7 Release an instance

Depending on the business change, you can manually release Pay-As-You-Go instances, but not Subscription instances. This document describes how to release an instance manually.

Attention

- Subscription instances are released automatically when they are overdue.
- The instance is in Running status.
- If the master instance enabled the read/write splitting function, to release the last read-only instance, 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.

4.8 Upgrade the database

Background information

RDS allows you to upgrade a database rather than downgrade it. See the console interface for upgradeable versions.

Attention

- Currently, this operation applies only to upgrades from MySQL 5.5 to MySQL 5.6 in the database.
- We recommend that you firstly purchase an instance with the database version you want to upgrade to and test its compatibility before upgrading.
- During the database upgrading 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 the low peak of the service, or 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.

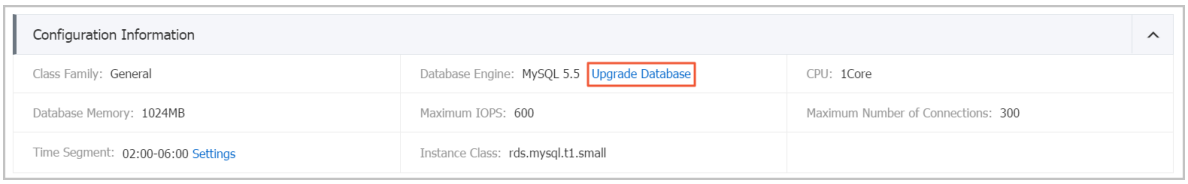

5. On the **Database Version Upgrade** page, select the target database version and click **Start Upgrade**.

4.9 Manually renew the subscription instance

A subscription instance must be renewed within 15 days of expiration. Subscription instances are auto-released when the payment is overdue for 15 days. After which, all the data for that instance is deleted and cannot be recovered. For more information on the 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 length on the **Renew** page.

> **Note:**
> You can choose to 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** and complete the payment process.

Related topic

Enable auto-renewal of the subscription instance

4.10 Enable auto-renewal of the subscription instance

Auto-renewal of the subscription instance frees you from regular manual renewals. It also avoids the service interruption caused when the instance expires if 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 primary 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 ##########.
- If you select auto-renewal option, 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, the 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 the next day, renew it manually to prevent service interruption.

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. Or 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, Select Automatic hours, as shown in the following figure.

7. Click Open automatic button.

4.11 Change configurations

Depending on your business needs, you can change the instance configuration, that is, change instance specifications, instance series (instance changed from base edition to available Edition), storage space, etc. During instance variation:

- During the change configuration take effect, the RDS service may have a 30-second flash, please do your best to perform the variation operation at the business low level, or make sure that your application has an automatic reconnection mechanism to avoid the impact of the burst
- The RDS supports setting the execution time for the variant operation.

Currently, only paid instances support the ability to change the configuration. This document describes how to change the configuration of RDS instances. For information about billing of configuration changes, see ######.

- Package year package month instance
  - In the same period of the contract, if you use real-time variation matching (the operation described in this article), the new configuration takes effect in real time, but can only be upgraded. If you choose to change the configuration when you are in, the new configuration takes effect at the start of the new billing cycle, supporting upfit or downfit.
  - After the instance expires, the configuration can be upgraded or degraded during the duration of the session, the new configuration takes effect at the start of the new billing cycle . For the maid step, see.
- Paid instances can be upgraded or degraded at any time.

Attention

- During configuration changes, you cannot perform most database, account, and network management operations. The following table lists the details. Choose a proper time to change the instance configuration.

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

Function	Supported or not
Switch to Subscription	No
Move 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 Public Address	No
Switchover between Primay and Slave Instances	No
Change Backup Mode	No
Restore Data	No
Modify Parameters	No

- Changing the configuration of an instance does not affect the data on the instance. However, when the system is applying the changes, a brief disconnection may occur (30 seconds). To minimize the impact, change the configuration when the service is not busy or make sure your application supports automatic reconnection.

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 variation instance page.
5. In the change configuration bar, select a new configuration, as shown in the following figure:

   Parameter description:

- Series: For switching between high-availability and financial edition, only available for instances of MySQL version 5.6.
- Availability area: You can choose to migrate an instance to another availability area, only available for MySQL 5.6 and SQL Server 2008 Release 2 instances.
- Specification: You can select an instance of other memory and CPU.
- Storage: select the appropriate storage space based on the usage of the current database storage space.

> **Note:**
> The storage space corresponding to each instance specification is different, if the storage space corresponding to the current specification does not meet your needs, first change the instance specification, then select space. For more information on specifications, see instance spec sheets.

- Switch time: select the execution time for the change instance configuration. Because changing the instance configuration involves bottom-level data migration, so you can choose to change the configuration 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 choose to perform an operation that changes the configuration of an instance in a maintable time.

6. If you want to modify maintenance time, do the following. If not, Skip.

   a. Click modify, as shown in the following figure, the system opens a new web page and jumps to the basic information page of the instance.

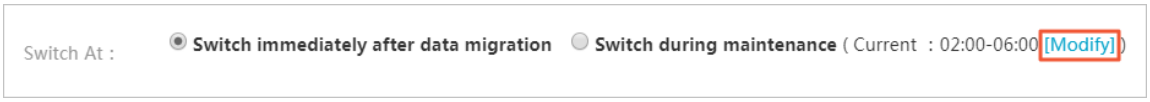

   b. In the configuration information bar, select the maintains time period, click Save, as shown in the following figure.

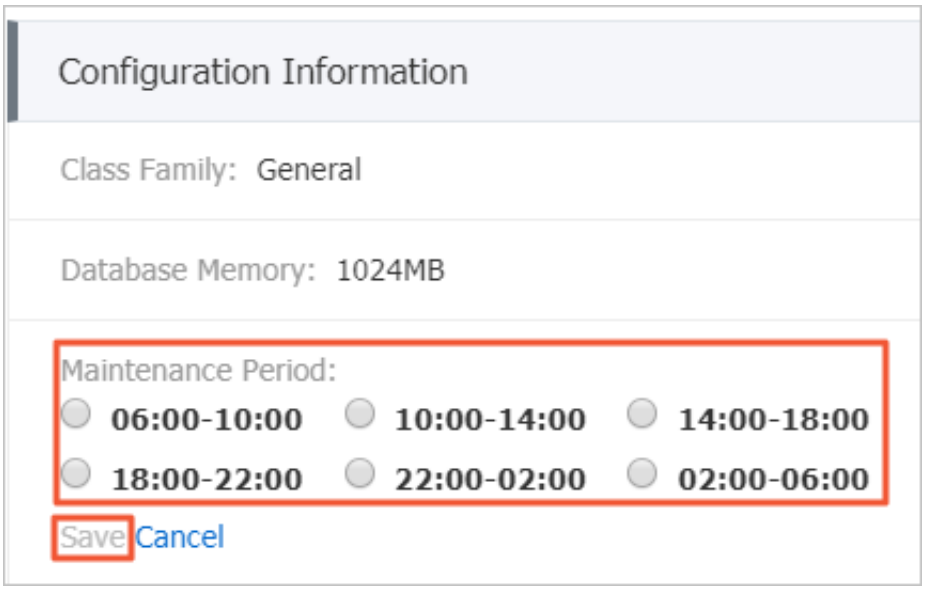

c. Returns the web page for a variant instance.

7. On the variation instance page, click confirm change, for years and months, please complete the payment process according to subsequent prompts.

4.12 SQL Server DBCC function

RDS SQL Server 2012 and later versions support some features of the 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 tracestatus(-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 command:

```
USE master
GO
--database engine edtion
SELECT SERVERPROPERTY('edition')
GO
```

```
--create database
CREATE DATABASE testdb
GO

DBCC tracestatus(-1)

exec sp_rds_dbcc_trace 1222,1

WAITFOR DELAY '00:00:10'

DBCC tracestatus(-1)
GO
```

4.13 End connection for SQL Server instances

Note:

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

Instances of RDS SQL Server 2012 and later versions are granted the End Connection (Kill) permission. However, you can only end the connection that you created and cannot end other connections, for example, backup connection.

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

4.14 Set instance parameters

4.14.1 Set parameters through RDS console

RDS allows you to define some instance parameters. For more information about the parameters that can be configured, see Parameter Settings on the RDS console. This document describes how to modify parameters and view the modification history on the RDS console. To perform these operations using APIs, see API references at the end of this article.

Note:

- PostgreSQL instances do not support user-defined parameters.
- To set parameters for instances of SQL Server 2012 and later versions, use SQL commands. For more information, see *Use SQL commands to set parameters*.

Background information

For descriptions about the database parameters, see the following official documents:

- *MySQL 5.5*
- *MySQL 5.6*

- *MySQL 5.7*
- *SQL Server*

Considerations

- Configure parameters only within the permissible value ranges shown on the parameter settings page.
- The instance must be restarted after modifying certain parameters. See the **Force Restart** parameter on the **Parameters** page to confirm if a restart is required. Before restarting, to avoid any interruption of production, you must guarantee the appropriate business arrangements. A restart will disconnect the instance. Exercise caution to restart the 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 instance to visit the **Basic Information** page.
4. Select **Parameters** on the left-side navigation pane.
5. Select the **Modifiable Parameters** tab.
6. Select the parameter modification method.

   - To modify a parameter

     1. Click the icon  following the parameter to modify.
     2. In the dialog box window, enter the target value in the field marked as 1 and click **OK**.
     3. Click **Submit modify** to confirm the setting, marked as 2 in the following figure.

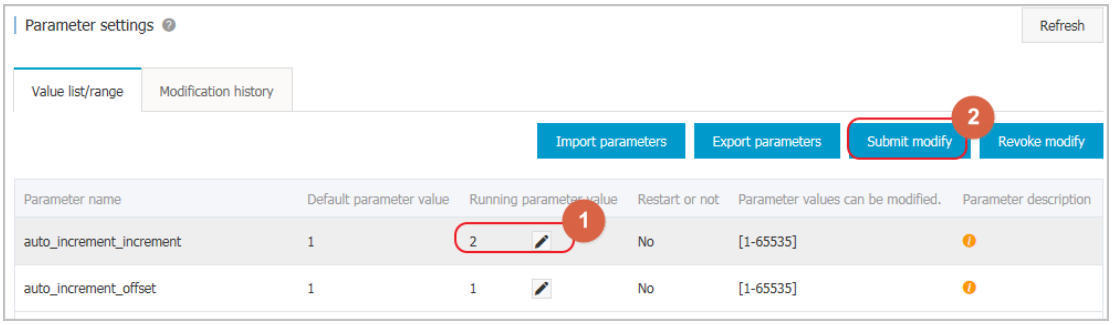

   - To modify multiple parameters:

     1. Click **Export Parameters** to export the parameter file (.txt) to your local device, marked as 1 in the following figure.
     2. Open the parameter file and batch modify the relevant parameters.

3. Click **Import Parameters**, marked as 2 in the following figure.
4. In the **Import Parameters** window, paste the parameters to modify and the parameter values and click **OK**, marked as 3 and 4 in the following figure.
5. Confirm the parameter modification results in the parameter list and click **Submit modify**, marked as 5 in the following figure.

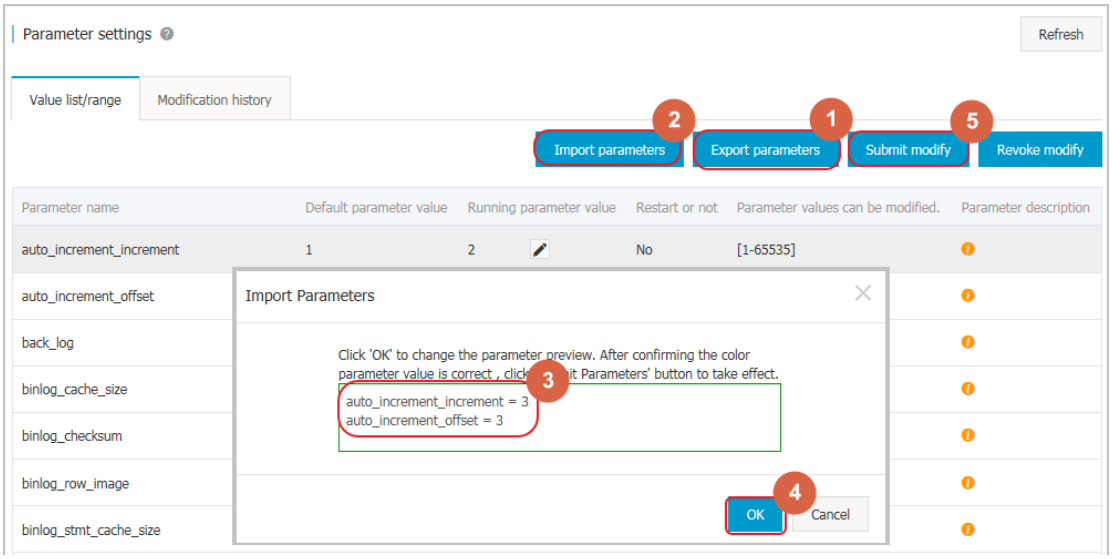

View the modification history

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. Select **Parameters** in the left-side navigation pane.
5. Select the **Modification History** tab.
6. Select the time range that you want to query, and click **Query**.

API references

- *DescribeParameterTemplates*
- *DescribeParameters*
- *ModifyParameter*

Best practices

Parameter optimization for MySQL instances

4.14.2 Use SQL commands to set parameters

Note:

The operation described in this document is applicable only to instances of RDS SQL Server 2012 and later versions. For the procedure of setting parameters for instances of other types and versions, see *Set parameters on the console*.

To set instance parameters, you only need to specify configuration options in the sp_rds_configure storage process. A prompt appears if the instance must be restarted to apply the parameter settings.

Currently, RDS only supports the following instance configurations:

- fill factor (%)
- max worker threads
- cost threshold for parallelism
- max degree of parallelism
- min server memory (MB)
- max server memory (MB)
- blocked process threshold (s)

Run the following command to set instance parameters:

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

SELECT *
FROM sys.configurations
WHERE NAME = 'max degree of parallelism'

EXEC sp_rds_configure 'max degree of parallelism',0

WAITFOR DELAY '00:00:10'

SELECT *
FROM sys.configurations
WHERE NAME = 'max degree of parallelism'
```

5 Account management

5.1 Create an account

You must create an account in the RDS instance before you can use the database. RDS supports two account modes: the classic mode and the master mode. The classic mode is an earlier management mode in which you cannot use SQL to manage databases and accounts. Master mode is a later management mode in which you can use SQL to manage databases and accounts. In addition, you have more permissions available in this mode. In the long run, master mode is recommended if you need personalized and fine-grained control over database permissions.

This document describes the features available for accounts in classic and master modes, and how to create accounts in different modes.

Account modes

In the classic mode, all accounts are created through the RDS console or API, instead of SQL. All accounts are created equal. The RDS console is used to create and manage all accounts and databases.

In the master mode, you must create and manage your first or initial account by using the RDS console or API. Then you can log on to a database using your initial account. When you are logged on, you can create and manage additional accounts using SQL commands. However, you cannot use your initial account to change the password for the additional accounts you have created. Instead, you have to delete those accounts and create new accounts. In the following example, the initial account is used as root to log on to the database. After that, an additional account “jeffrey” is created:

```
mysql -hxxxxxxxxx.mysql.rds.aliyuncs.com -uroot -pxxxxxx -e "
CREATE USER 'jeffrey'@'%' IDENTIFIED BY 'password';
CREATE DATABASE DB001;
"
```

In master mode, the database management page is unavailable on the RDS console for now. APIs such as *CreateDatabase* cannot be used to manage databases. Instead, you must use SQL commands or DMS to create and manage databases.

The following figure shows how to create and manage databases or accounts in classic and master modes:

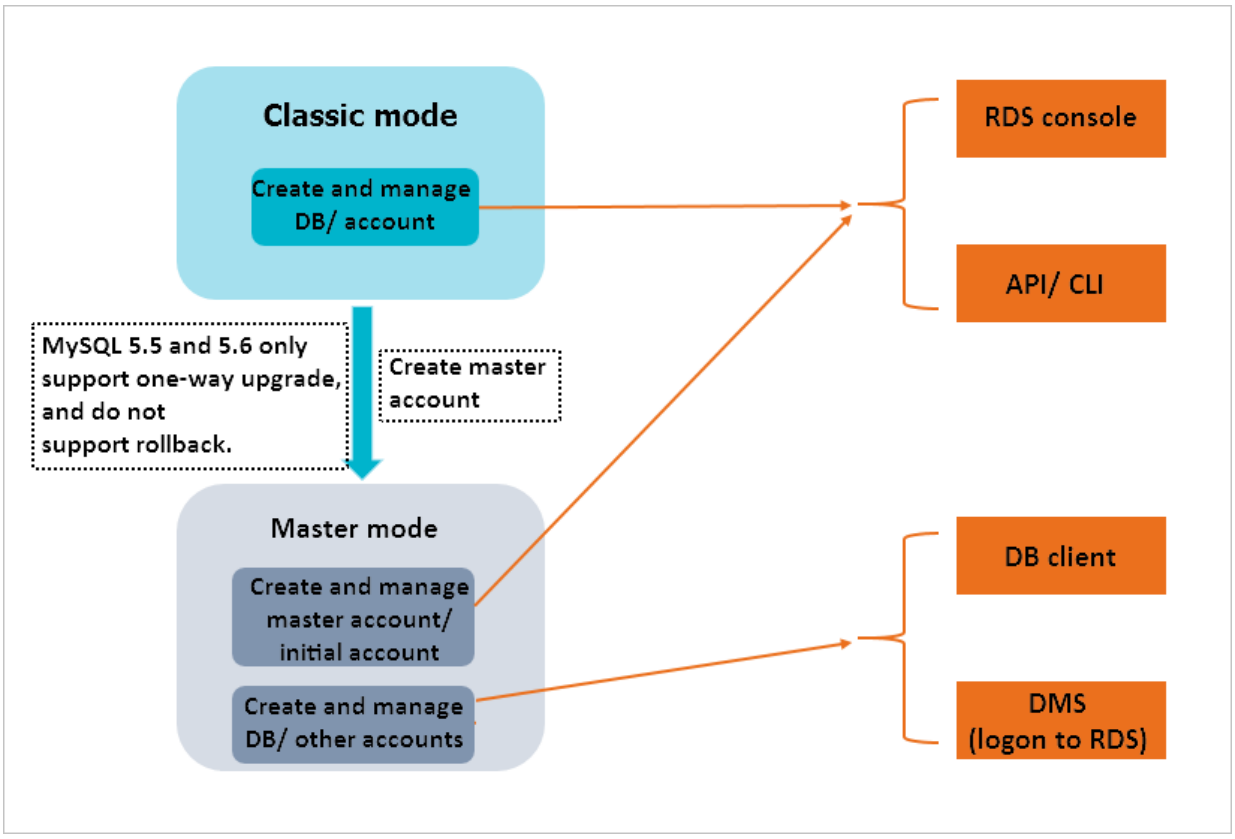

Comparison between the classic and master modes

Account modes available for database engines

The account modes available for various database engines are shown as follows:

Database engine	Account mode
MySQL 5.5/5.6	Classic mode/Master mode **Note:** Upgrade from classic to master mode is supported only. You cannot roll back after the upgrade.
MySQL 5.7	Master mode
SQL Server 2008 R2	Classic mode
SQL Server 2012/2016	Master mode
PostgreSQL	Master mode
PPAS	Master mode

Differences between accounts and permissions

The following table lists the differences between classic and master modes in accounts and permissions:

Item	Classic mode	Master mode
Account limit	Up to 500.	No limit.
Database limit	• MySQL: Up to 500. • SQL Server: Up to 50.	No limit.
RDS console used to create and manage databases and accounts	Yes	• The console can be used to manage the first account created on it, but not additional accounts, which must be created and managed using SQL commands or DMS. • Instead of the console, SQL commands or DMS must be used to create and manage databases.
SQL used to manage databases and accounts	No	Yes
Permission management	Simple: Choose from Read/Write or Read-Only permissions for each account.	Fine-grained control. You can take full advantage of the database engine's permission management capabilities. For example, you can assign the query permissions for different tables to different users.
Permissions for an account (applicable to MySQL only)	SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, PROCESS, INDEX, ALTER, CREATE TEMPORARY TABLES, LOCK TABLES, EXECUTE, REPLICATION SLAVE, REPLICATION CLIENT, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, EVENT, TRIGGER	CREATE USER, RELOAD, and REFERENCES are supported in addition to the 20 permissions supported in classic mode.

Difference in features

There are no differences in product features in both classic and master modes, including read-only instances, read/write splitting, configuration upgrade, network management, IP address whitelisting, and monitoring and alarms.

How to create an account

Attention

- When assigning database account permissions, follow the minimum permissions principle and service roles to create accounts and assign reasonable Read-Only and Read/Write permissions. When necessary, you may split database accounts and databases into smaller units so that each database account only has access to its own service data. If you do not must write data to a database, please assign Read-Only permission.
- Use strong passwords for database accounts and change the passwords on a regular basis.

Procedure

- See the following documents for more information about how to create an account in classic mode:
  - *MySQL 5.7####/5.5/5.6########*
  - *########SQL Server 2008 R2#*
- See the following documents for more information about how to create an account in master mode:
  - *#######*
  - *MySQL 5.7############*
  - *########*
  - *########*

5.2 Reset 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 to change the password periodically.

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, select the account you want to reset the password for and click **Reset Password**.

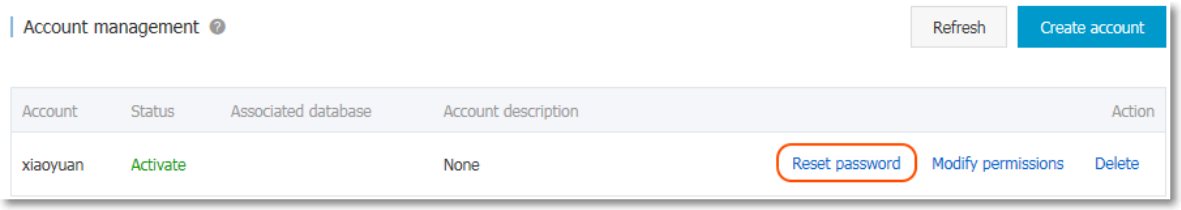

4. On the **Reset Account Password** page, enter a new password and click **OK**. The password consists of 6 to 32 characters, which must be letters, digits, hyphen (-), or underscores (_). We do not recommend using your old passwords.

5.3 Change account permissions

While using RDS, you can change permissions of the account any time as per 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.

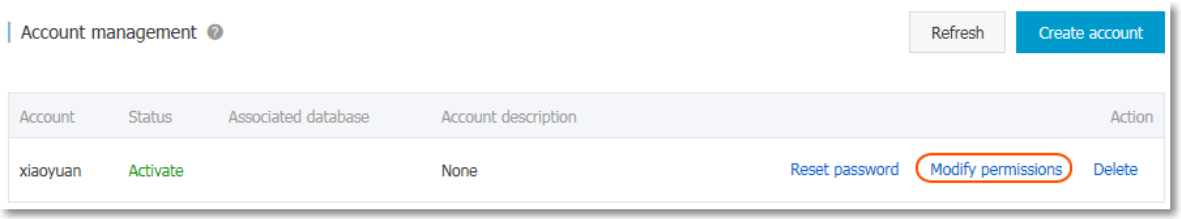

4. On the **Modify account** page, change the account permissions and click **OK**, as shown in the following figure.
   - Add an authorized database: Select a database in **Unauthorized database** and then click **Authorize** to add it to **`Authorized database`**.
   - Delete an authorized database: Select a database in **Authorized database** and then click **<** Remove to add it to **`Unauthorized 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`** , you can also click **Grant All Read/Write** or **Grant All Read-only**. (Note: Either of them is displayed at a time.)

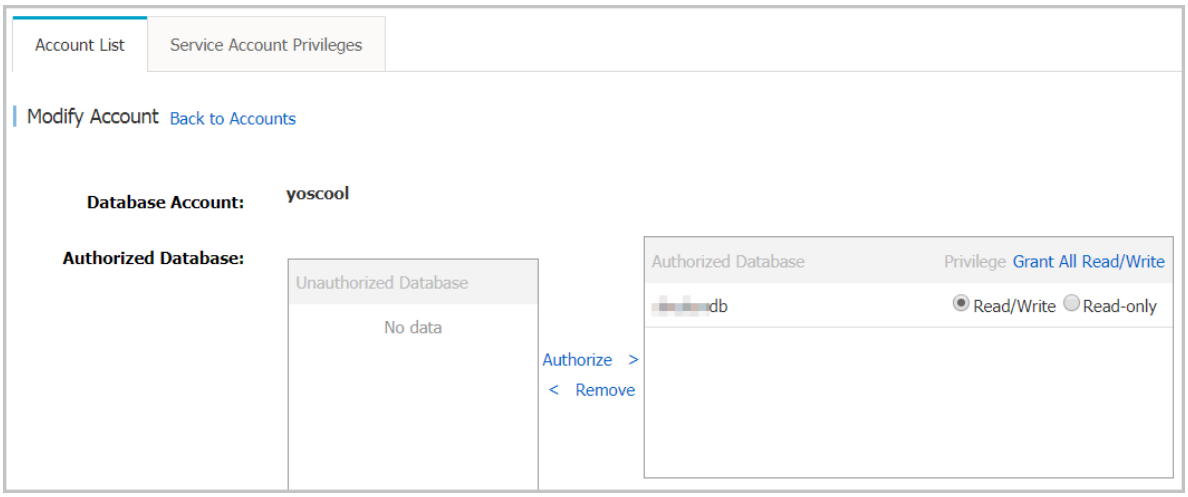

5.4 Authorize a service account

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 configuration or view table structure , index, and SQL, 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 white lists, 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**.

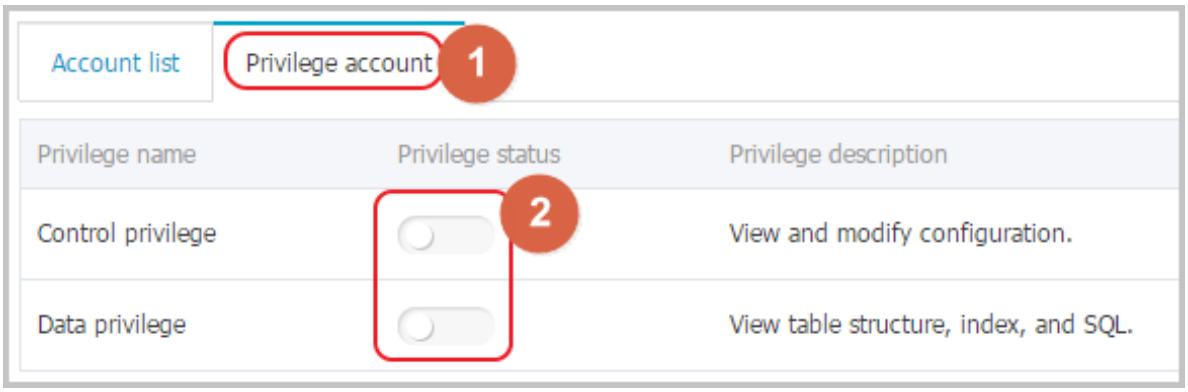

5. After setting the permission expiration time on the page of **Setting expired time**, click **OK** as shown in the following figure.

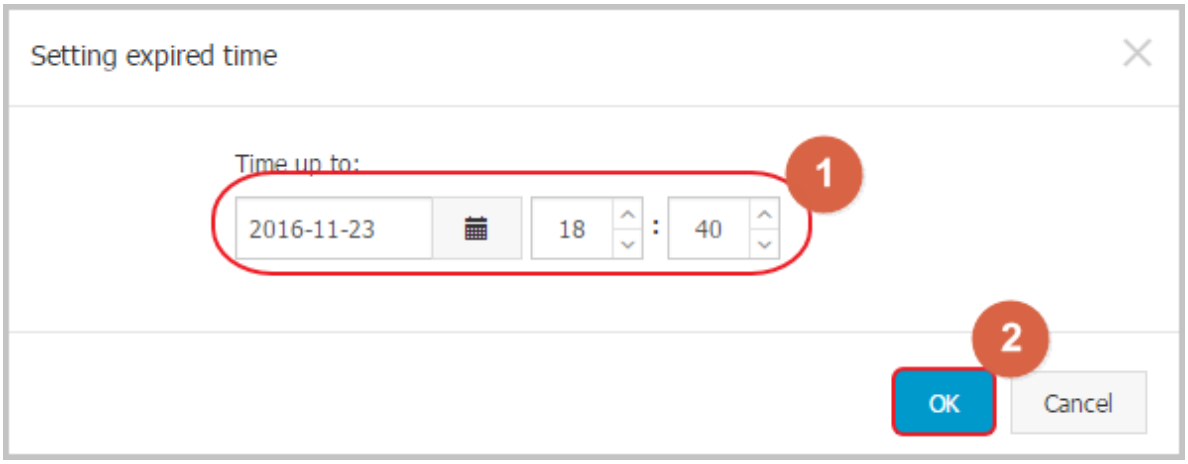

Subsequent operations

After a service account is authorized, you may cancel the authorization (as shown in Fig.1) or change the authorization validity period (as shown in Fig.2) on the **Privilege account** tab page.

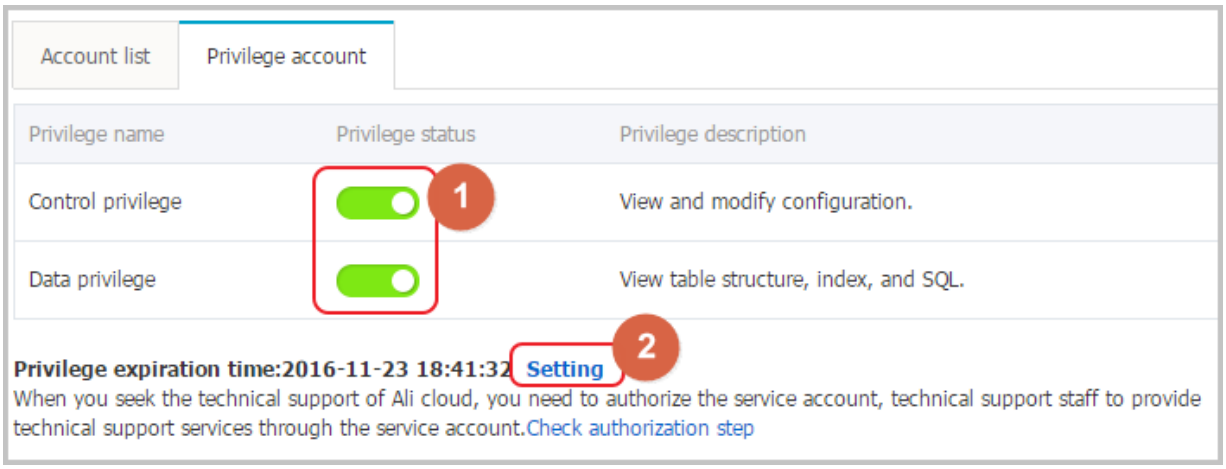

5.5 Delete an account

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

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 corresponding 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 accounts 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';
```

5.6 LOGIN user management of SQL Server instances

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

> **Note:**

> The operation described in this document is applicable only to instances of RDS 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-8325465CA9B7'
```

When creation is in progress, the LOGIN user is assigned permissions at the server level and database level. The Message area shows the following information.

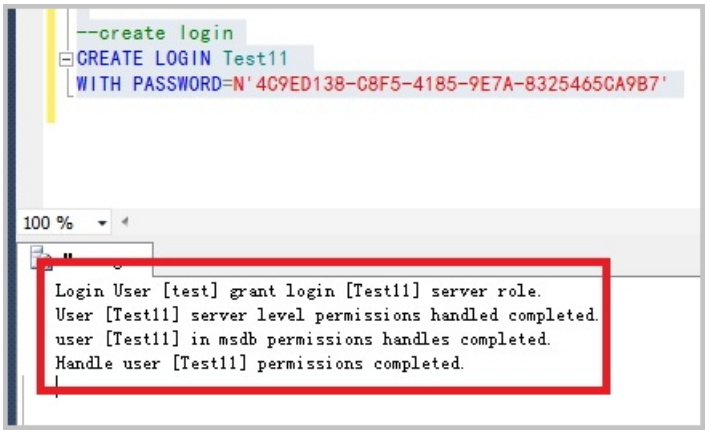

Modify LOGIN

Run the following command to modify a LOGIN user.

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

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

5.7 User management of SQL Server instances

You can create common users in the user database that you created, but not in the system database. This document describes how to create and manage users in a database of ApsaraDB for SQL Server by using SQL commands.

Note:
The operation described in this document is applicable only to instances of RDS 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 command 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 instructions of SQL Server. For example, you can run the following command to modify user-mapped logon informatio n:

```
USE TestDB
GO
```

```
ALTER USER test WITH LOGIN=test
```

Delete a user

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

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

6 Read/write splitting

6.1 Introduction to read/write splitting

Functions

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

- The connection address of the master instance: The Internet and intranet address can coexist.
- The connection address of the read-only instance: The Internet and intranet address can coexist.
- The read/write splitting address: The Internet and intranet address cannot coexist. An intranet address is generated by default. If you must use the read/write splitting Internet address, you can switch the 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 the application to split read and write operations.

With this function, an extra read/write splitting address is provided to link the master instance with all its read-only instances, which enables an automatic link for read/write splitting. 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 handling capacity of the system by adding more read-only instances without making any changes to the application.

The following figure shows the different connection addresses.

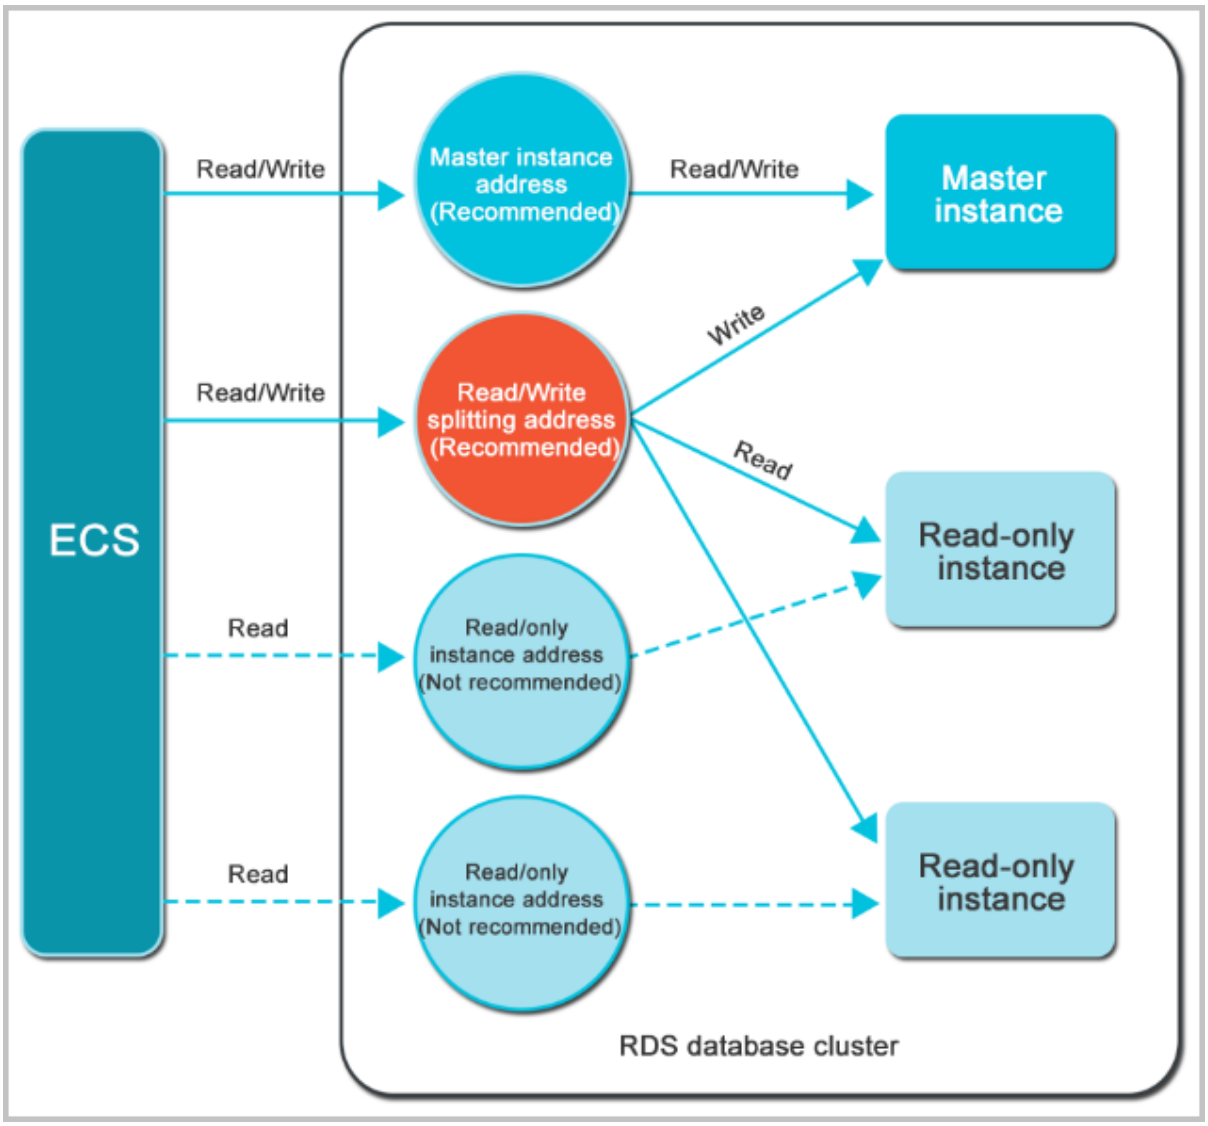

Benefits

- 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 the applications so that write requests go 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/write operations on the master and read-only instances, with the read/write requests automatically distributed. Therefore, maintenance costs are reduced.
- Improves performance with RDS 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, impacting response latency significantly. RDS read/write splitting can be directly set in the existing highly secure link without time consumption by any other components, which reduces latency and improves processing speed.

- Applies to various use cases 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 check.

  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, the application still runs normally when any single-node read-only instance fails. After the instance resumes, RDS reclaims it automatically 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 service.

  RDS provides all read-only instance users with the read/write splitting function for free.

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 `stmt close`.
  - The environment configuration variables of `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 processlist`, `show master status`, and `com_process_info` commands return results according to the instance connected during execution.
- All transactions are routed to the primary 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 primary database or encapsulate query requests in transactions.
- The `LAST_INSERT_ID()` function is not supported. To use this function, add `hint :/* FORCE_MASTER*/, eg:/*FORCE_MASTER*/ 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*.
- The database proxy is enabled for the instance. To enable the database proxy, see *Database proxy*.

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.

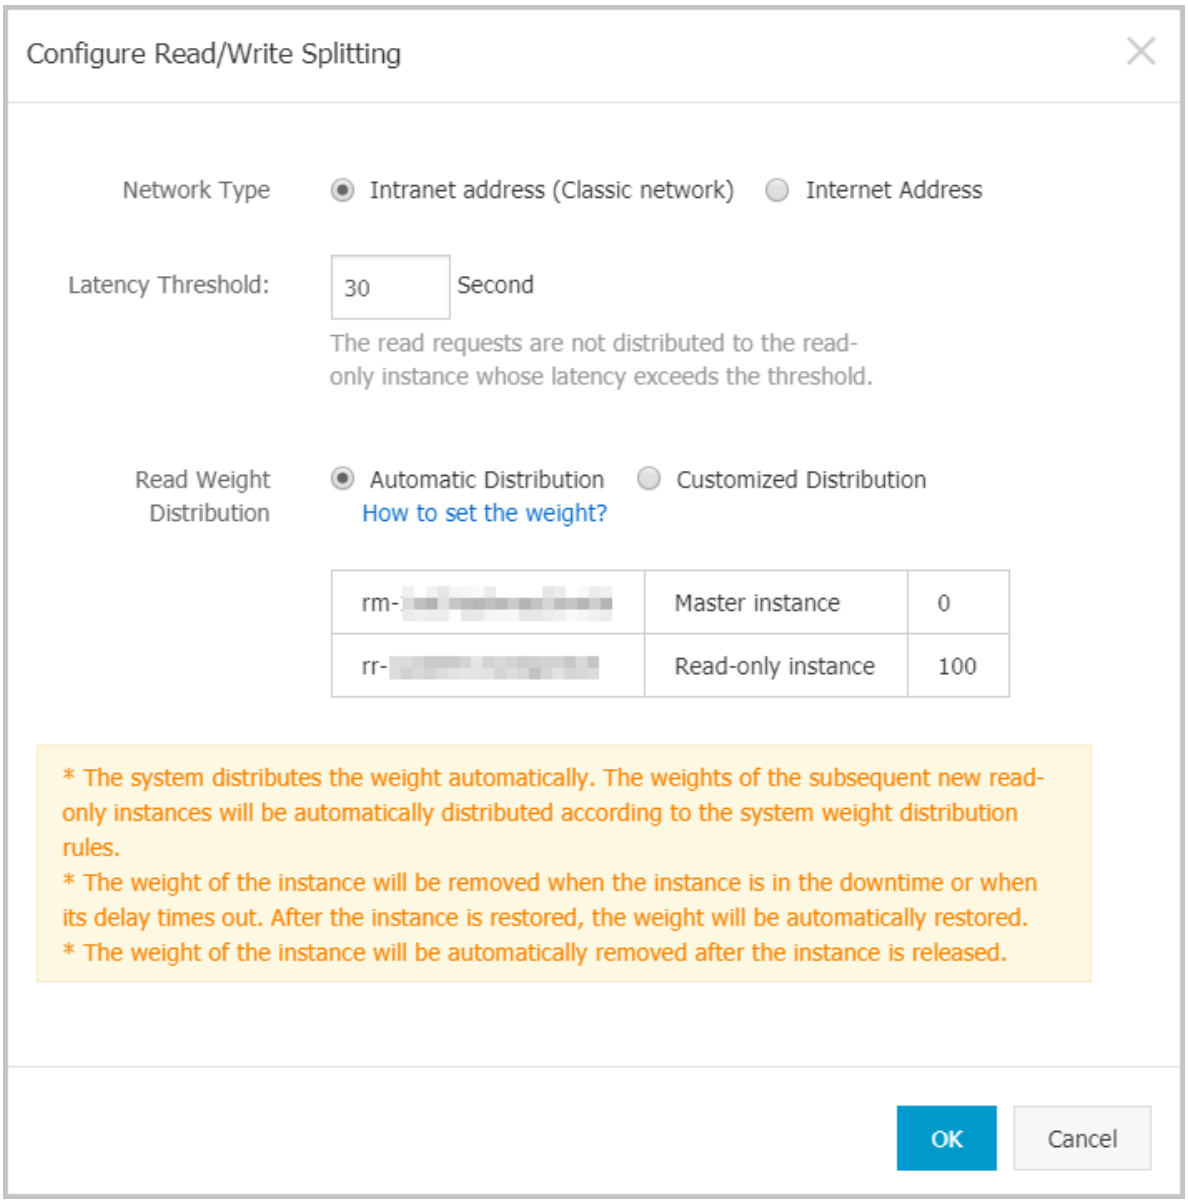

Parameter descriptions:

- 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:** After you click **OK**, the instance status changes to **Creating Network Connection**. Wait for a while patiently. After the instance status changes to **Running**, enter the **Read/Write Splitting** page. After the read/write splitting function is successfully enabled, the interface appears as shown in the following figure. The following figure is for reference only.

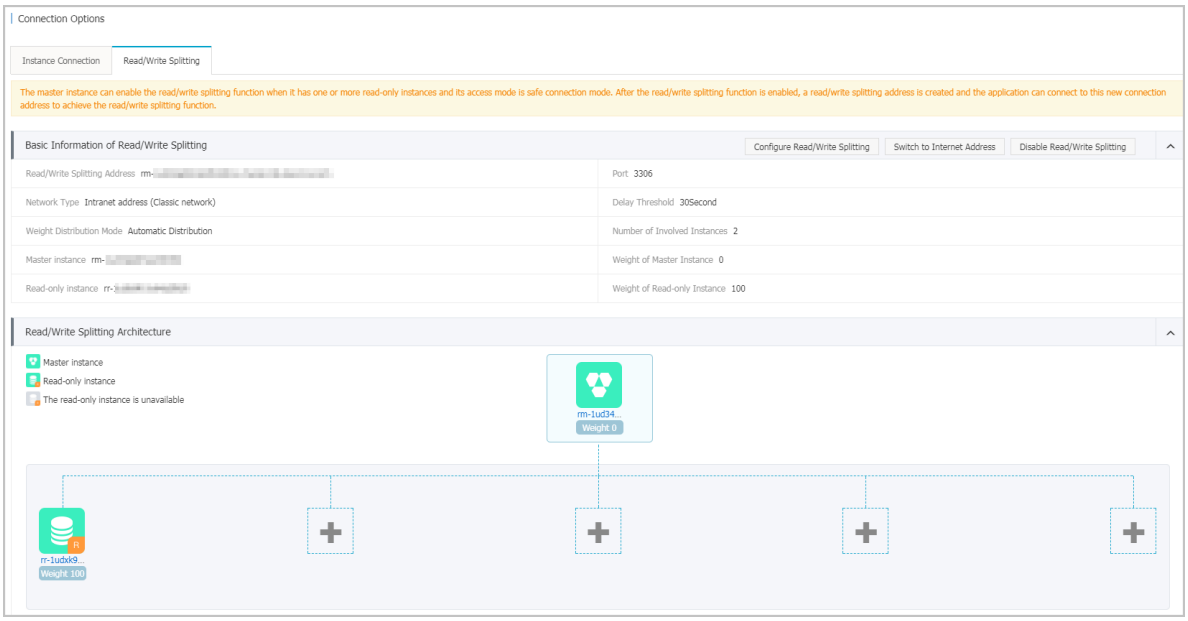

6.3 Modify the latency threshold and read weight distribution

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. Select **Connection Options** in the left-side navigation pane to enter the **Connection Options** page.
5. Select **Read/Write Splitting** tab.
6. Click **Configure Read/Write Splitting** to enter the **Configure Read/Write Splitting** page.

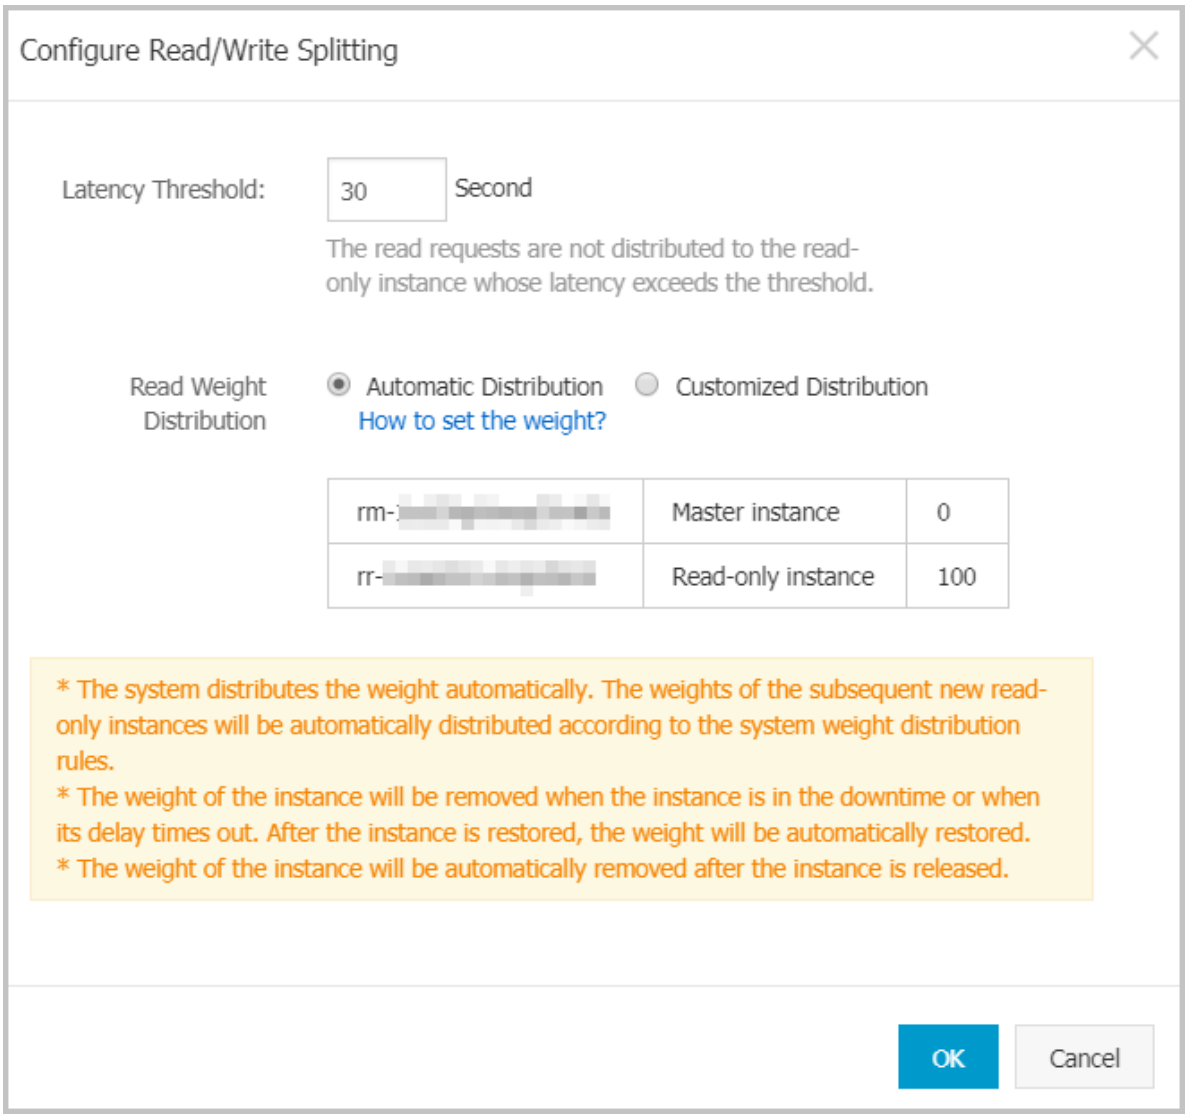

7. Change settings as needed by referring to the following information:

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

- Latency Threshold: This refers to the latency threshold of read-only instances and ranges from 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 SQL queries, latencies may occur in read-only instances. We recommend that you set the value to at least 30s.
- Read Weight Distribution: This refers to the read request weights of different instances. An instance with a higher weight processes more read requests. For example, if the 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 the weights, use either of the following modes:

- Automatic Distribution: The system automatically distributes read weights to instances (including read-only instances added afterwards) according to their specifications. For more information about the read weights of instances with different specifications, see *Rules of weight distribution by system*.
- Customized Distribution: You can set the read weights of instances with a value ranging from 0 to 10,000. In this mode, the weights of newly added read-only instances are 0 by default, and must be changed manually.

> **Note:**
> To enable certain queires to return data in a real-time manner, you can forcibly forward these statements to the master instance using hints. 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 switch the read/write splitting address type based on business scenarios. When you enable the read/write splitting function, the read/write splitting intranet address is generated by default. This article introduces how to switch between the intranet and Internet addresses of read/write splitting.

Prerequisites

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

Precautions

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

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 **Switch to Internet Address**.

   Note:
   Note: If you are switching from the Internet address to intranet address, click **Switch to Intranet Address**.
7. On 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 can delete the last read-only instance. Otherwise, the deletion does not succeed.

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

Note:
After the read/write splitting function is disabled, your application 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 a 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. Select **Connection Options** in the left-side navigation pane to enter the **Connection Options** page.
5. Select the **Read/Write Splitting** tab.

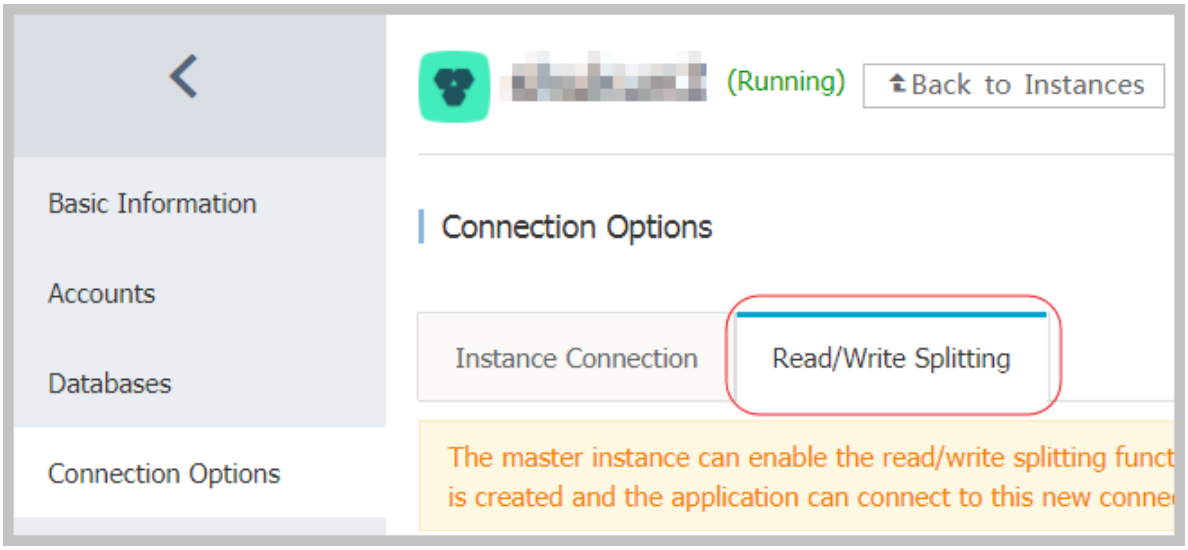

6. Click **Disable Read/Write Splitting**.
7. In the dialog box, click **Confirm**.

6.6 Monitor read/write splitting performance

You can view the read/write splitting performance by using 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. Select **Monitoring and Alarms** in the left-side navigation pane to enter the **Monitoring and Alarms** page.
5. Select the monitoring type **Engine Monitoring** in the **Monitoring** tab, as shown in the following figure.

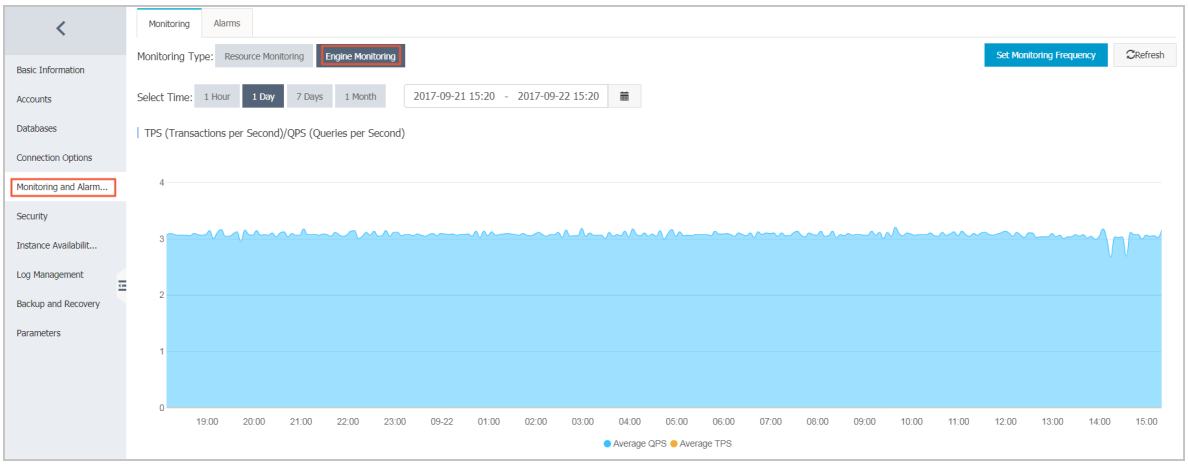

6. You can see the number of reads/writes of each database (master database and read-only database involved in the read/write splitting) by viewing the TPS (Transaction Per Second)/ QPS (Query Per Second) data.

6.7 Test read/write splitting performance

After the 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 article describes how to correctly configure its parameters to conduct read/write splitting performance testing.

Prerequisites

- The read/write splitting function is enabled. See *Enable read/write splitting* for detailed procedures.
- The Sysbench 0.5 is installed. See *Sysbench documentation* for instructions on downloading and installing Sysbench 0.5.

Attentions

- We recommend that a case with prepare or a transaction not be used to test the load balancing of the read/write splitting function.
- Avoid the master/slave latency from exceeding the threshold set for the monitoring check due to high stress on write requests.
- We recommend that you use the following Sysbench scripts to build a specific SQL statement as needed.

```
function thread_init(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 benchmark the read/write splitting, 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 test, 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	Indicates whether the data is randomly initialize d.
max-time	Stress testing duration.
max-requests	Total number of requests in a stress testing duration.
num-threads	Number of concurrent threads.
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 benchmark tests as needed.

Name	Description
oltp-test-mode	Test mode. But this parameter is unavailable in the Sysbench 0.5, so ignore this parameter. Possible values: • complex: Transactional test as the default value. • Simple: Simple test for read-only SQL statements. • nontrx: Non-transactional test. • sp: Stored procedures.
oltp-skip-trx	Indicates whether begin at the beginning of the SQL statement and commit at the end of the SQL statement 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 benchmark the read/write splitting performance, you must set it as “on” and omit the begin/commit at the beginning and end of the SQL statement.
oltp-read-only	Indicates whether read-only SQLs 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 parameter values as needed to perform read-only or read/write tests.

Stress testing examples

Test read/write performance

1. Run prepare command.

```
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 run command.

> **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` is not effective, it indicates that 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 cleanup command.

```
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 prepare command.

```
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 run command.

```
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 cleanup command.

```
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/write splitting effect

6.8.1 Use SQL audit to verify the read/write splitting effect

You can use SQL audit logs to verify the effect of read/write splitting by comparing the number of SQL statements executed by the master instance and that of the read-only instances that participate in read/write splitting.

To enable SQL audit and view SQL audit logs, see *#unique_73*.

6.8.2 Use internal SQL commands to verify the read/write splitting effect

You can verify the effect of read/write splitting by running the `/*PROXY_INTERNAL*/show last route;` command.

Note:
Do not use this SQL statement in the production environment now, because it is still being tested internally and may be adjusted later as required.

View the database to which an SQL command is sent for execution

You can view the ID of the instance on which an SQL command runs by running the following SQL command:

```
/*PROXY_INTERNAL*/show last route;
```

Note:
RDS provides a built-in hint SQL command (which can only be executed using read/write splitting VIP). When running this command on the MySQL client, always append the -c option to the command, or the hit is filtered out by the client and the following error is returned.

```
ERROR 1064 (42000): You have an error in your SQL syntax; check the
manual that
```

```
corresponds to your MySQL server version for the right syntax to use
near 'last route' at
line 1
```

The returned result last_bkid indicates the ID of the database to which the last SQL command (the one before the hit) was sent. This ID is the unique identification of a RDS instance and is unique to the instance. The following figure shows the details.

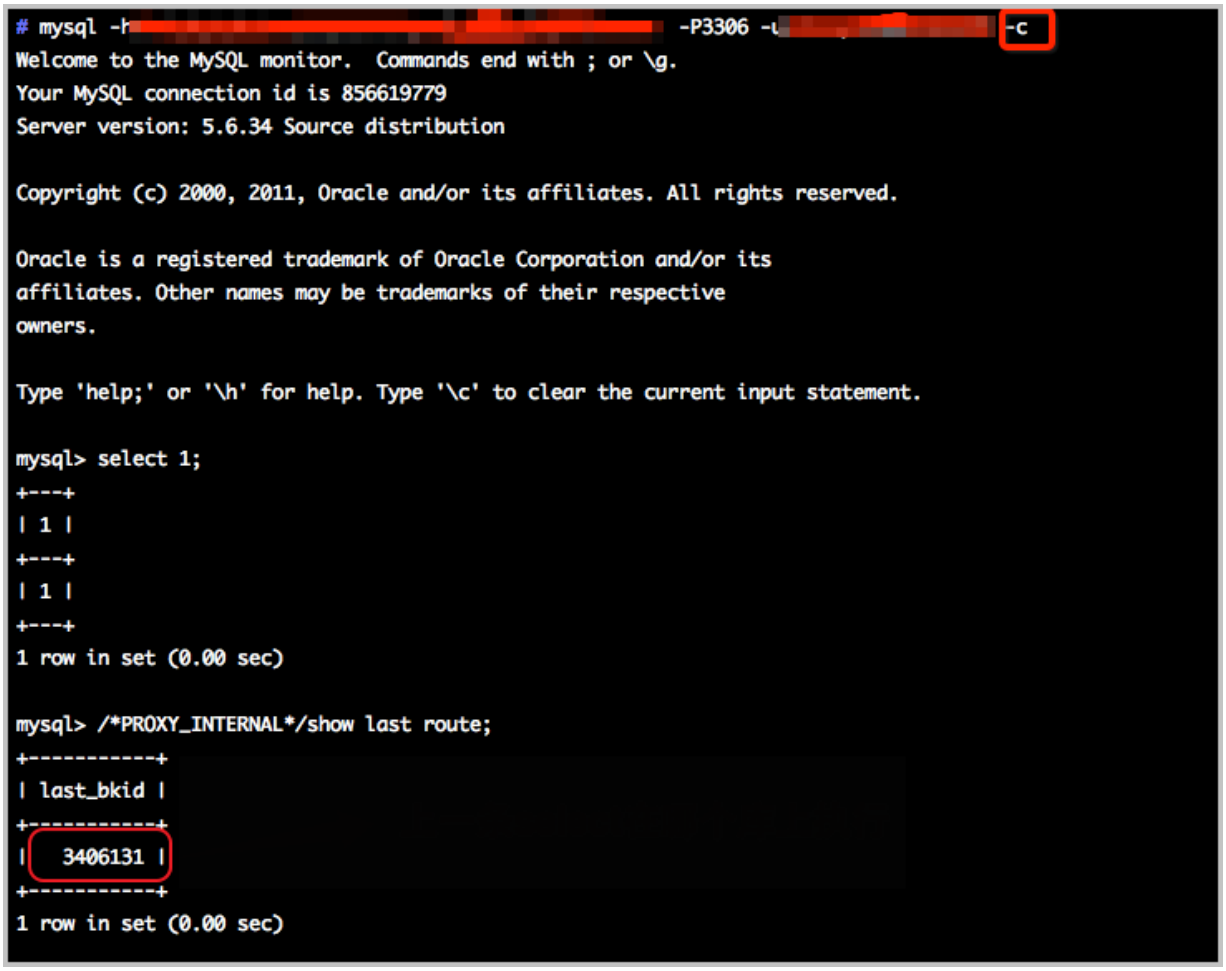

Note:

Given that the SQL load for RDS is measured in batch with a minimum unit of 100 entries, the first 100 select statements are executed on the same instance ID before the 101st select statement is routed to another instance ID. To verify this, you can write a simple SQL file such as the following a.sql file:

```
select 1;
/*PROXY_INTERNAL*/show last route;select 1;
***100 entries***;
select 1;
```

```
/*PROXY_INTERNAL*/show last route;
```

Now you can see that the 101st SQL statement is routed to another instance ID (assuming that more than two read-only instance IDs are available for supporting the load).

Verify that all write requests are forwarded to the master database (master instance) for execution

Once the read/write splitting is enabled for RDS instances, write requests can only be forwarded to the master database because read-only databases only process read requests. Even when a system or routing error occurs, for example a write SQL statement is routed to a read-only database, such write request can be routed back to the master database for execution according to the error reason (read_only error).

Besides, you can run the insert statement and then the following hint SQL statement to verify that all write requests are forwarded to the master database.

```
/*PROXY_INTERNAL*/show last route;
```

Verify that all read requests are forwarded to slave databases (read-only instances) for execution

Query the ID of the instance that executes read requests by running the following hint SQL command to verify that these read requests are forwarded to a slave database.

```
/*PROXY_INTERNAL*/show last route;
```

Note:

Given that the SQL load for RDS is measured in batch with a minimum unit of 100 entries, the first 100 select statements are executed on the same instance ID before the 101st select statement is routed to another instance ID. To verify this, you can write a simple SQL file such as the following a.sql file:

```
select 1;
/*PROXY_INTERNAL*/show last route;select 1;
***100 entries***;
select 1;
/*PROXY_INTERNAL*/show last route;
```

Now you can see that the 101st SQL statement is routed to another instance ID (assuming that more than two read-only instance IDs are available for supporting the load).

6.9 Verify read weight distribution

To verify the load ratio of read weight, you can run the `select @@server_id;` command for 10,000 times using persistent connections and count 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 with the 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. Select **Monitoring and Alarms** in the left-side navigation pane to enter the **Monitoring and Alarms** page.
5. Select the monitoring type **Engine Monitoring** in the **Monitoring** tab.
6. You can see the number of reads/writes of each database (master database and read-only database involved in the read/write splitting) by viewing the TPS (Transaction Per Second)/ QPS (Query Per Second) data.

   Note:
   The refresh of 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 each database

You can view the number of SQL statements run by each instance by connecting the master database and read-only databases involved in read/write splitting.

Note:
To perform this verification, the connection addresses of the master database and read-only databases instead of the read/write splitting addresses are needed.

Run any 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;
```

6.10 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 Instance	240 MB	3	100
rds.mys2.mid	Common Instance	600 MB	5	100
rds.mys2.standard	Common Instance	1,200 MB	6	400
rds.mys2.large	Common Instance	2,400 MB	9	400
rds.mys2.xlarge	Common Instance	6,000 MB	10	800
rds.mys2.2xlarge	Common Instance	12,000 MB	10	800
rds.mys2.4xlarge	Common Instance	24,000 MB	12	1000
rds.mys2.8xlarge	Common Instance	48,000 MB	13	1000
rds.mysql.t1.small	Common Instance	1 GB	1	100
rds.mysql.s1.small	Common Instance	2 GB	1	100
rds.mysql.s2.large	Common Instance	4 GB	2	200

Specification code	Specification type	Memory	CPU	Weight
rds.mysql.s2.xlarge	Common Instance	8 GB	2	200
rds.mysql.s3.large	Common Instance	8 GB	4	400
rds.mysql.m1.medium	Common Instance	16 GB	4	400
rds.mysql.c1.large	Common Instance	16 GB	8	800
rds.mysql.c1.xlarge	Common Instance	32 GB	8	800
rds.mysql.c2.xlarge	Common Instance	64 GB	16	1600
rds.mysql.c2.xlp2	Common Instance	96 GB	16	1600
rds.mysql.c2.2xlarge	Common Instance	128 GB	16	1600
mysql.x8.medium.2	Dedicated Instance	16 GB	2	200
mysql.x8.large.2	Dedicated Instance	32 GB	4	400
mysql.x8.xlarge.2	Dedicated Instance	64 GB	8	800
mysql.x8.2xlarge.2	Dedicated Instance	128 GB	16	1600
rds.mysql.st.d13	Dedicated Host	220 GB	30	3000
rds.mysql.st.h13	Dedicated Host	470 GB	60	6000

Specify whether an SQL is sent to the master instance or a read-only instance by using Hint

In addition to the weight distribution system of read/write splitting, Hint serves as a complement ary SQL syntax to specify whether an SQL is executed on the master instance or a read-only instance.

The Hint formats supported by RDS read/write splitting are as follows:

- `/*FORCE_MASTER*/`: specifies that the SQL is executed on the master instance.
- `/*FORCE_SLAVE*/`: specifies that the SQL is executed on a read-only instance.

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

```
/*FORCE_MASTER*/ SELECT * FROM table_name;
```

7 Database management

7.1 Create a database

You can create databases using the *RDS console*. Database names are unique within an instance, but can be duplicate across instances.

- For MySQL 5.7 Basic Edition, see *MySQL 5.7############* to create a database using a client.
- For SQL Server 2012 or 2016, see *########SQL Server 2012#* to create a database using a client.

Background information

- Databases under a single instance share all the resources of this instance.
  - A MySQL instance supports up to 500 databases.
  - A SQL Server 2008 R2 instance supports up to 50 databases.
  - A PostgreSQL or PPAS instance has no limit on the number of databases.

Procedure

1. Log on to the *RDS console* and select the target instance.
2. Select **Database Management** in the left-side navigation pane, and click **Create database**, as shown in the following figure.

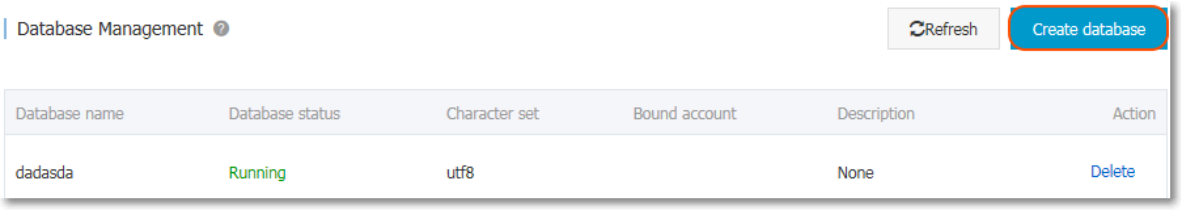

3. Enter the information of the database you want to create, and click **OK**, as shown in the following figure.

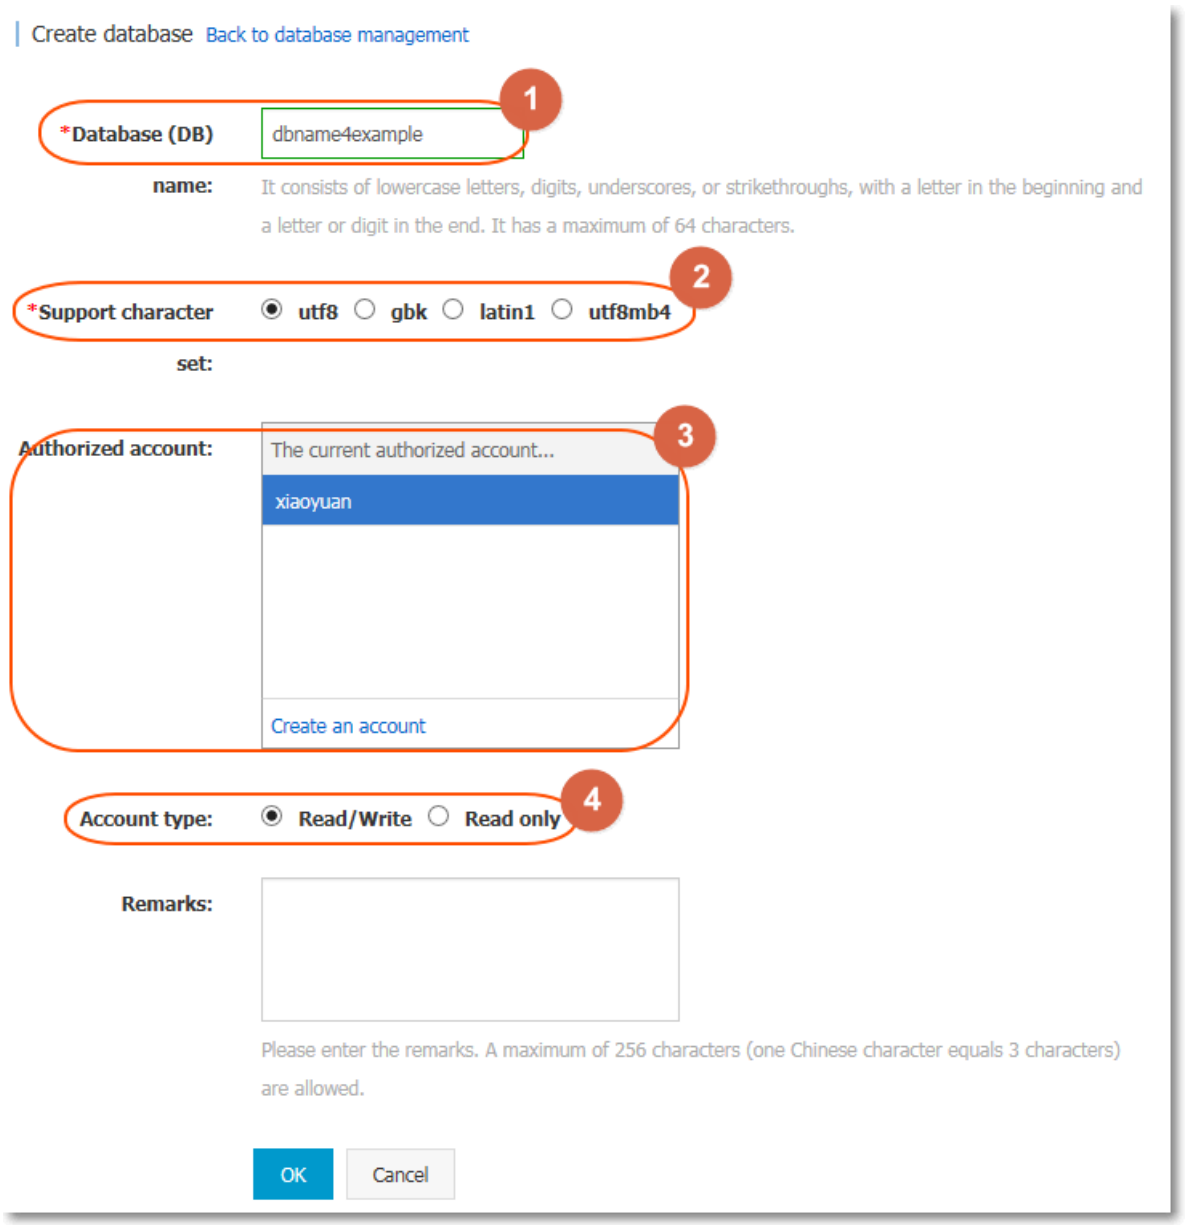

- Database (DB) name: Contains 2 to 64 characters, which consist of lowercase letters, digits , underscores (_), or hyphens (-). It must begin with a letter and end with a letter or digit.
- Support character set: utf8, gbk, latin1, and utf8mb4.
- Authorized account: Select an account authorized by this database. This field can be blank if no account has been created.
- Account type: This option is visible after **Authorized Account** is selected. Set the permission authorized by this database to **Authorized Account**, which can be set to **Read/ Write** or **Read-only**.
- Remarks: This field can be used to store additional information relevant to the database to facilitate management. A maximum of 256 characters can be entered.

7.2 Delete a database

This section describes how to delete a database using SQL statements or the RDS console.

Currently, you can delete databases of SQL Server 2008 R2, MySQL 5.5, MySQL 5.6 instances on the RDS console.

Note:
You can only use SQL statements to delete a database if a master account has been created for MySQL 5.5 or 5.6 instances.

Delete a database from 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 instance to visit the **Basic Information** page.
4. Select **Database Management** in the left-side navigation pane.
5. Locate the database you want to delete and click **Delete** in the **Action** column.
6. In the displayed confirmation dialog box, click **OK**.

Delete a database using SQL statements

Currently, you can delete databases of MySQL 5.7, PostgreSQL, SQL Server 2012, and PPAS using SQL statements.

1. To connect to an RDS instance through a client, see:
   - ####
   - ####
   - ####
   - ####
2. Run the following command to delete a database:

   ```
   drop database <database name>;
   ```

7.3 Copy a database for SQL Server

7.3.1 Copy a database for SQL Server 2008 R2

You can copy an existing database to produce an identical one. This article describes how to copy a database on the RDS console and applies only to SQL Server 2008 R2 instances. For instances

of SQL Server 2012 and later versions, you can copy a database only by using SQL commands. For more information, see *Copy a database for SQL Server 2012 or later versions*.

Attention

- Only one database can be copied at a time.
- The new database must be named differently from all the existing databases.

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, select **Databases** to go to the **Database** page.
5. Click **Copy Database**.
6. Set the parameters.

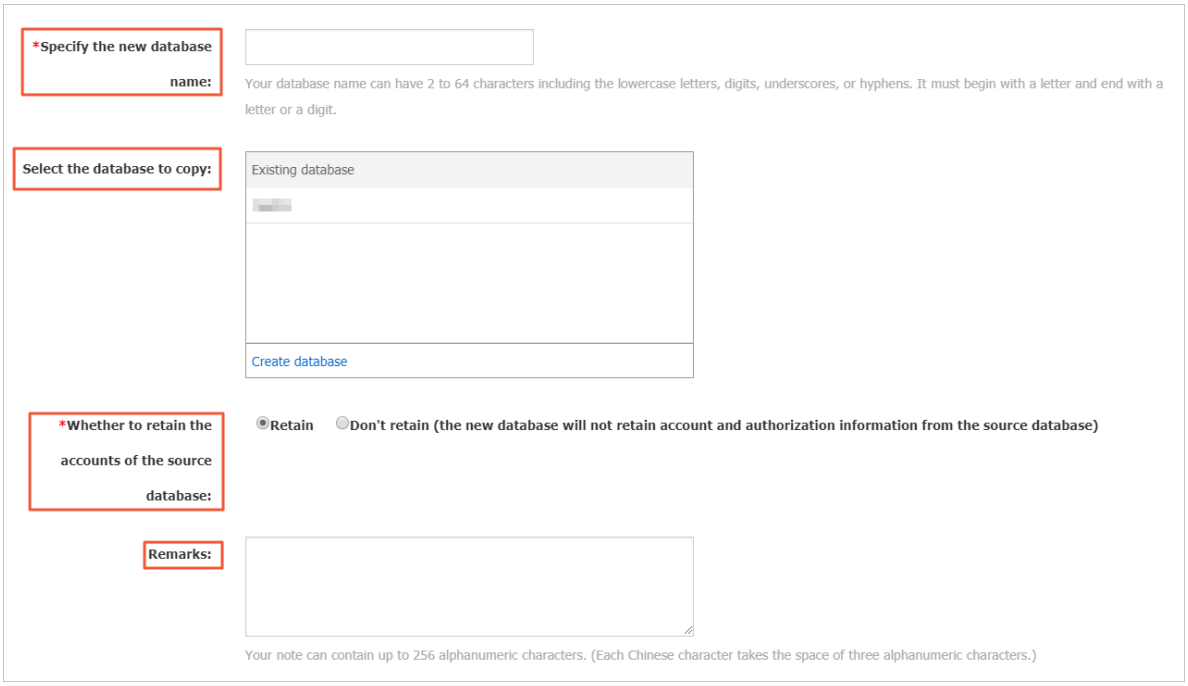

Parameter description:

Parameter name	Description
Specify the new database name	The new database name consists of up to 64 characters including lowercase letters, digits, underscores (_), and hyphens (-). It must start with a letter and end with a letter or digit.

Parameter name	Description
Select the database to copy	Select the database that you want to copy from the list of existing databases.
Whether to retain the accounts of the source database	When copying a database, you can choose whether to transfer the account and permissions from the source database to the new database. The default option is **Retain**, which means transferring the accounts.
Remarks	You can add information about the new database to facilitate subsequent database management. A maximum of 256 characters can be entered.

7. Click **OK**.

7.3.2 Copy a database for SQL Server 2012 or later versions

Note:

This article is applicable only to instances of SQL Server 2012 and later versions. For information about how to copy a database for SQL Server 2008 R2, see *Copy a database for SQL Server 2008 R2*.

To clone a database by using SQL commands, specify the source and target databases in the stored procedure sp_rds_copy_database. The cloning duration varies depending on the size of the database.

Prerequisite

Before cloning a database, make sure that the available space of the instance is at least 1.3 times the size of the cloned database.

Procedure

Run the following commands to clone a database:

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

GO
EXEC sp_rds_copy_database 'testdb','testdb_copy'
```

```
SELECT *
FROM sys.databases
WHERE name IN ('testdb','testdb_copy')

SELECT
    family_guid,database_guid,*
FROM sys.database_recovery_status
WHERE
DB_NAME(database_id) IN ('testdb','testdb_copy')
```

7.4 Database management of SQL Server instances

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

Note:
The operation described in this document is applicable only to instances of RDS 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 while running the following command:

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

  The following error message is 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 transaction 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 while running the following command:

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

  The following error message is 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 transaction ended in the trigger. The batch has been aborted.
  ```

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

  For example, if you run the following ONLINE command directly:

  ```
  USE [master]
    GO

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

    ALTER DATABASE [TestDb]
    SET ONLINE
  ```

  The following error message is 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] .
```

8 Network management

8.1 Set access mode

This function has been replaced by the database proxy function. For more information, see *Database proxy*.

8.2 Set network type

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

> **Note:**
> To migrate an instance from a classic network to a VPC without stopping services, see *Hybrid access solution for the seamless migration from classic network to VPC*.

Background information

On Alibaba Cloud platform, a classic network and a VPC have the following differences:

- Classic network: The cloud services in a classic network are not isolated, and unauthorized access can be blocked only by the security group or whitelist policy of the 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 migrate applications to the cloud seamlessly.

Precautions

- After switching the network type, the original intranet IP address is changed and the public 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 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 from a classic network to a VPC, the access mode must be safe connection mode. To switch the access mode, see *Set access 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 the network type switching, the RDS service may be interrupted for about 30 seconds . Therefore, perform the switching 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. Select **Connection Options** in the left-side navigation pane to open the **Connection Options** page.
5. Do the following to switch the network type:
   - Switch from a classic network to a VPC
     1. Click **Switch to VPC**.
     2. Select a VPC and virtual switch.

        > **Note:**
        > - If the drop-down lists do not have 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 *Set access mode*.

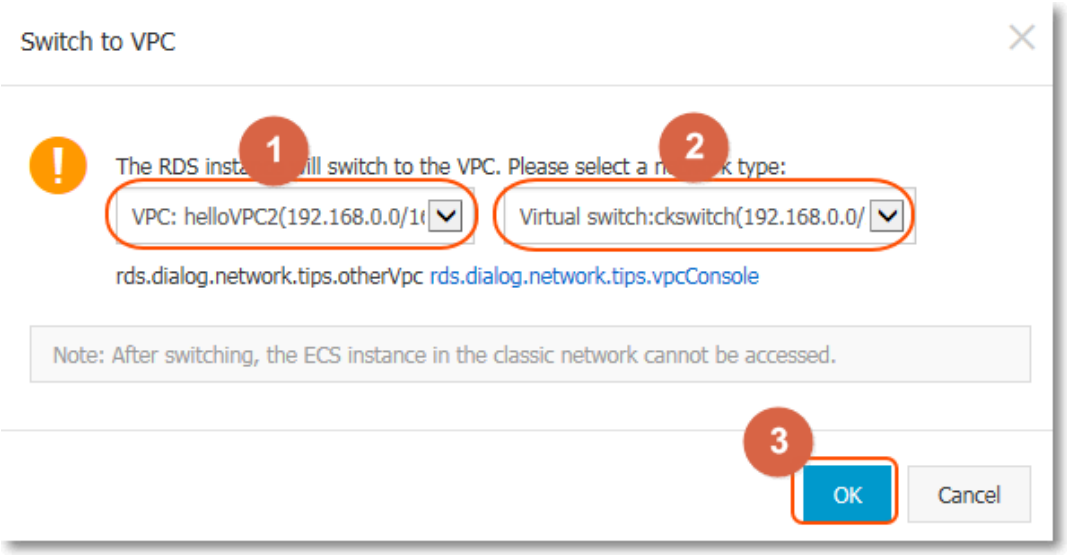

3. Click **OK**.

- Switch from a VPC to a classic network

1. Click **Switch to Classic Network**.
2. Click **OK**.

8.3 Hybrid access solution for the seamless migration from classic network to VPC

Virtual Private Cloud (VPC) is a private network logically isolated from other virtual networks. Alibaba Cloud VPC allows you to build an isolated network environment with better security and performance than classic network. With these benefits, VPC has 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 seamless migration from classic network to VPC 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 VPC without any interference with each other.

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

Background information

With a traditional solution, migrating an RDS instance from classic network to VPC causes immediate release of the intranet address of the classic network. That would result in an intermitte nt interruption for up to 30 seconds and prevent the ECS on the classic network from accessing

the RDS instance by using the intranet 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 from classic network to 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 seamless migration, RDS has added the hybrid access feature, which makes 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 classic network and VPC. During the hybrid access period, the RDS instance reserves the intranet address of the original classic network and adds an intranet address under VPC, which prevents any intermittent interruption during migration. We recommend that you use VPC only for the sake of security and performance. For this reason, hybrid access is available for a limited period of time. That means the intranet address of the original classic network is released when the hybrid access period expires. In that case, your applications cannot access the database using the intranet address of the classic network. You must configure the intranet address under VPC in all your applications during the hybrid access period, so as to guarantee the seamless network migration and minimize the impact on your services.

For example, a company wants to migrate from classic network to VPC. They can use the hybrid access solution to have a transitional period in which some of their applications can access the database by VPC, and the others can continue to access the database using the intranet address of the original classic network. When all the applications can access the database by VPC, the intranet address of the original classic network can be released, as shown in the following figure.

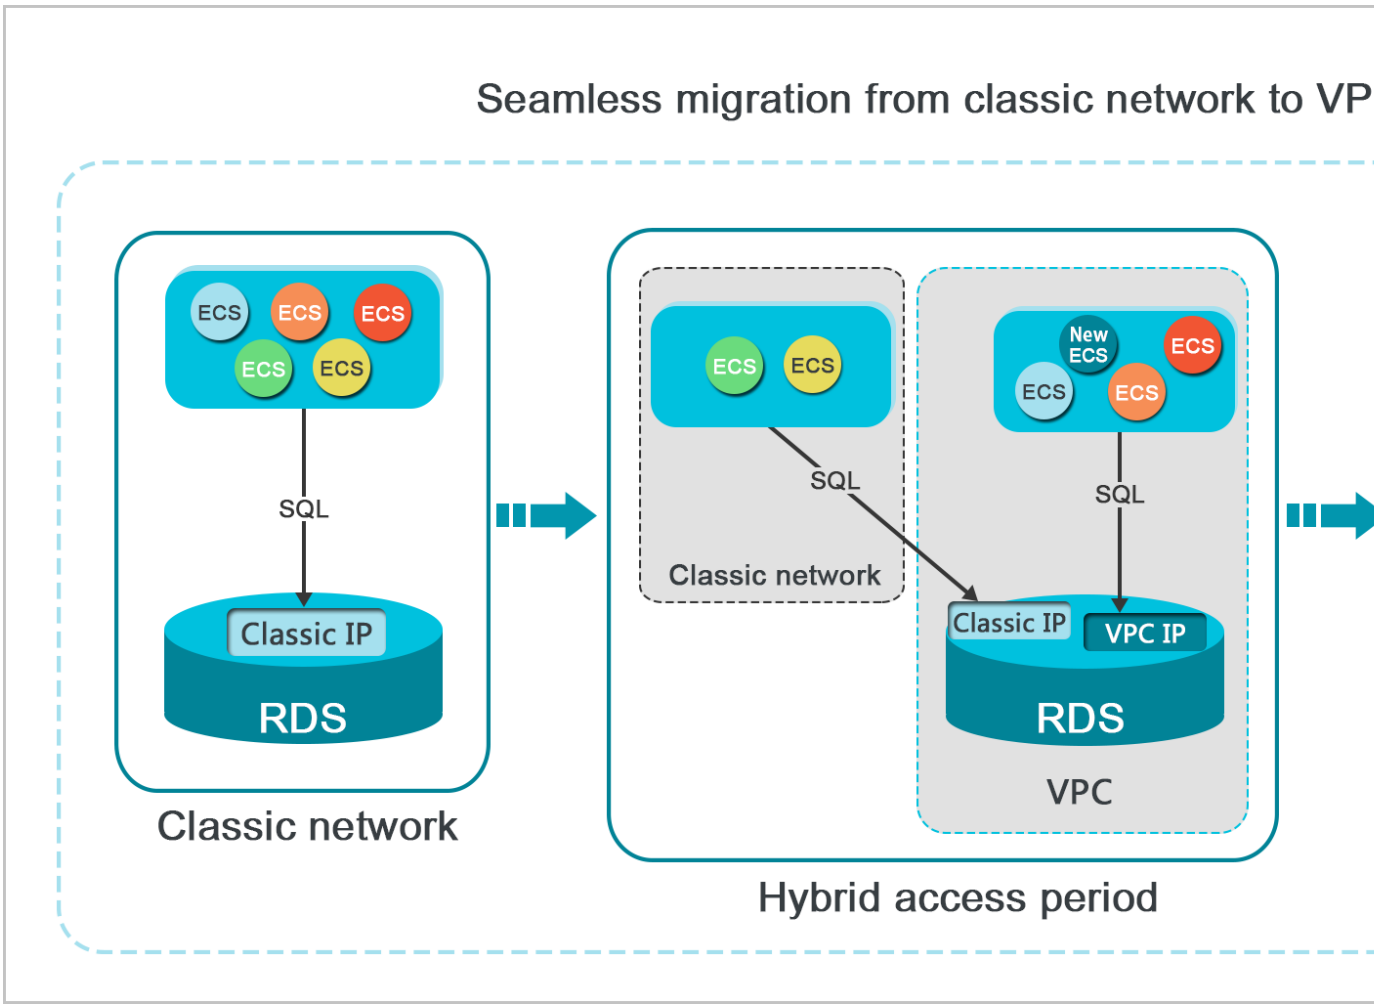

Functional Limits

The following functional limits apply during the hybrid access period:

- Switch to classic network is not supported.
- Zone migration is not supported.
- Switch between the High-availability Edition and Finance Edition is not supported.

Prerequisites

- The current access mode is safe connection mode. For more information on how to switch the access mode, see *Set access mode*. MySQL 5.7, SQL Server 2012 and SQL Server 2016 only support standard mode, but these instances also support hybrid access in this condition.
- The current network type is classic network.
- There are available VPC and VSwitch in the zone where the RDS instance is located. If not, see *Create VPC* and *Create VSwitch* to create one.

Migration 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, select **Connection Options** to enter the **Connection Options** page.
5. On the **Instance Connection** tab, click **Switch to VPC**.
6. On the **Switch to VPC** confirmation page, select the VPC and Vswitch to switch to.
7. Check **Reserve original classic endpoint**, and select the **Expiration time** for the basic intranet address of the original network, as shown in the following figure.

Note:

- Starting from the seventh day before the date on which the intranet address of the original classic network is set to be released, the system sends a text message of a notice to the mobile number listed on your account every day.
- When the reservation ages out, the intranet address of the classic network is released automatically and can no longer be used to access the database. To prevent service interruption, set the reservation period as necessary. After the hybrid access configuration is complete, you can change the expiration date.

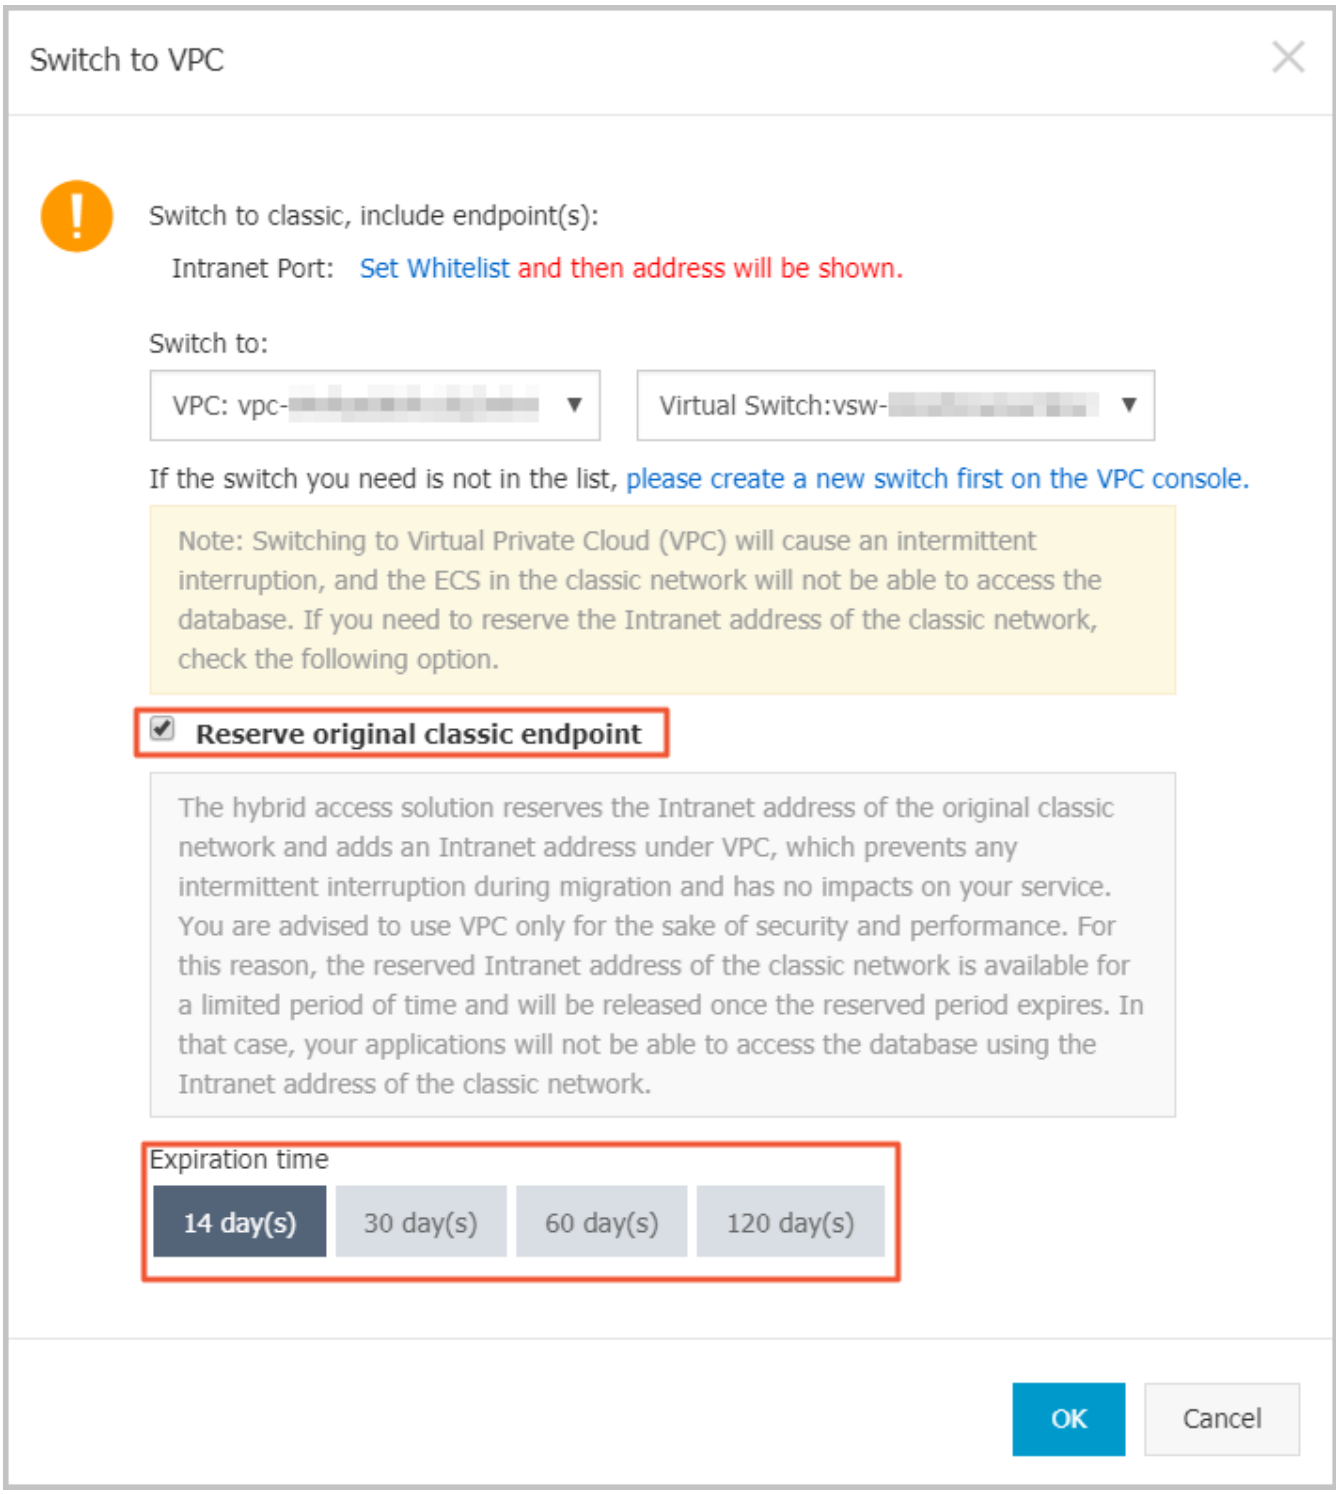

8. Click **OK**, and the **Original classic endpoint** area appears on the console, as shown in the following figure.

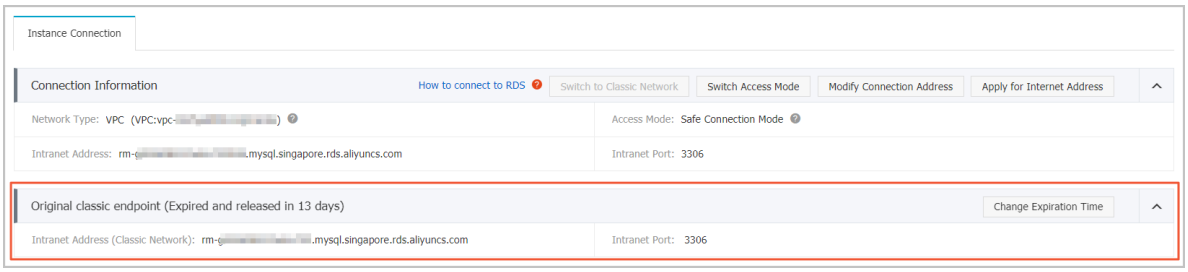

Change the expiration time of the original classic endpoint

During the hybrid access period, you can change the reservation period of the intranet 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 address of the original classic network is set to expire on 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, select **Connection Options** to enter the **Connection Options** page.
5. On the **Instance Connection** tab, 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**.

8.4 Set intranet and Internet addresses

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

Background information

RDS supports connections through the intranet addresses and Internet addresses. 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 address • Internet address • intranet and Internet addresses
High-availability Edition	• MySQL 5.5/5.6 • SQL Server 2008 R2 • PostgreSQL 9.4 • PPAS 9.3	Standard mode	• Intranet address • Internet address
		Safe connection mode	• Intranet address • Internet address • intranet and Internet addresses
Finance Edition	MySQL 5.6	Standard mode	• Intranet address • Internet address
		Safe connection mode	• Intranet address • Internet address • intranet and Internet addresses

The applicable scenarios of the connection addresses are as follows:

- Use the intranet address only:
  - The system provides an intranet 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 with those of your RDS instance.
- Use the Internet address only:
  - This scenario is applicable when your application is deployed on the ECS instance that is located in the different region with that of your RDS instance.
  - This scenario is applicable when your application is deployed on the platform other than Alibaba Cloud.
- Use both of the intranet and Internet 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 on the ECS instances of different regions at the same time.
- 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 the platform other than Alibaba Cloud at the same time.

Attentions

- Before accessing the database, you must add the IP addresses or IP segments used to access the database to a whitelist. For more information, see Set whitelist.
- The RDS charges a fee for traffic over the Internet address. For detailed charges, see *RDS Pricing*.
- Connecting to the RDS instance with an Internet address reduces the instance security. To get a higher transmission rate and a higher security level, we recommend that you migrate your applications to the ECS instances in the same region with that of your RDS.

Apply for the Internet 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. Select **`Connection options`** in the left-side navigation pane.
5. Click **Apply for Internet Address**, as shown in the following picture.

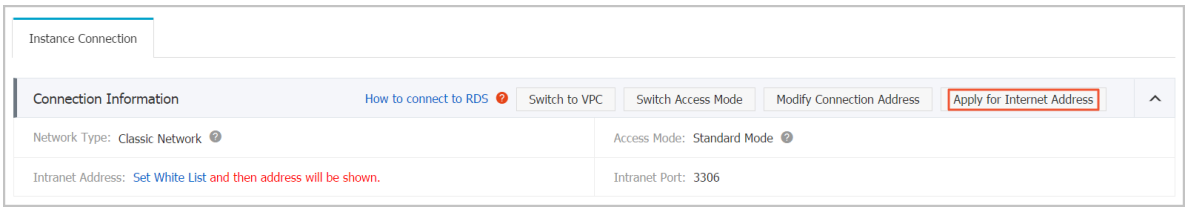

6. On the displayed confirmation window, click **OK** to generate an Internet address.

Modify the connection address

You can modify the Internet and intranet connection address as per your needs. Do as follows:

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. Select **Connection options** in the left-side navigation pane.
5. Select the **Instance Connection** tab.
6. In the **Connection Information** 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.

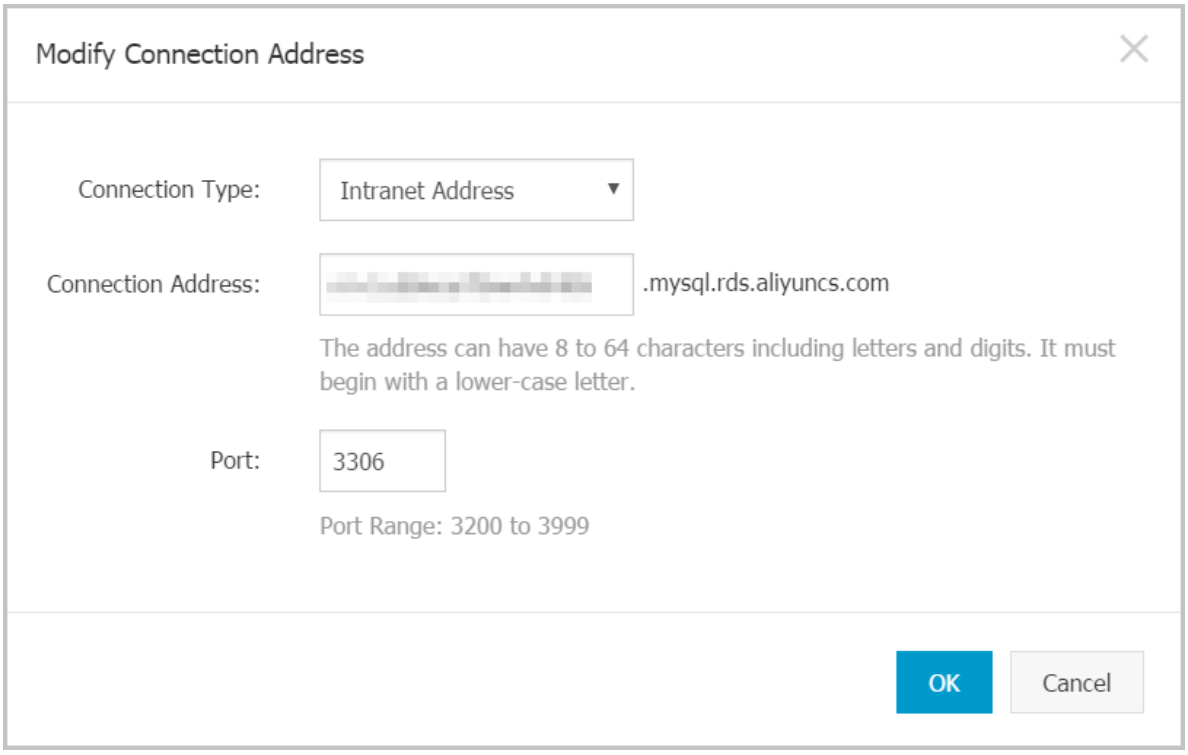

Parameters description:

- 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 can have 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 the Internet address

If you must release the Internet address, do as follows:

Note:

The operation is displayed under the **safe connection mode**. For more information about safe connection mode, see *Set access 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. Select **Connection options** in the left-side navigation pane.
5. Select the **Instance Connection** tab.
6. In the **Connection Information** area, click **Release Internet Address**.

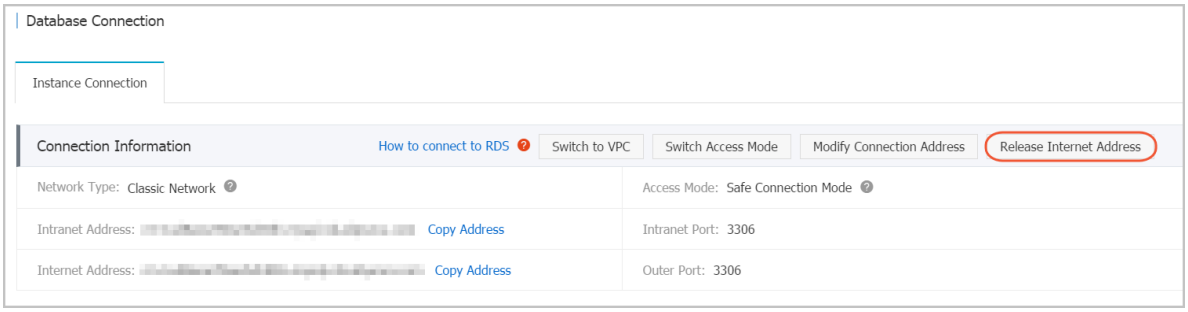

7. Click **Confirm** on the displayed confirmation interface to release the Internet address.

9 Security management

9.1 SQL audit

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

Attention

- Certain RDS instance types do not support the SQL audit function.
- The SQL audit function does not affect the 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 the analysis based on network protocols. 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, SQL audit logs are not accurate incremental data.
- 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 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, select **Security**.

5. Select the **SQL Audit** tab and click **Enable now**.

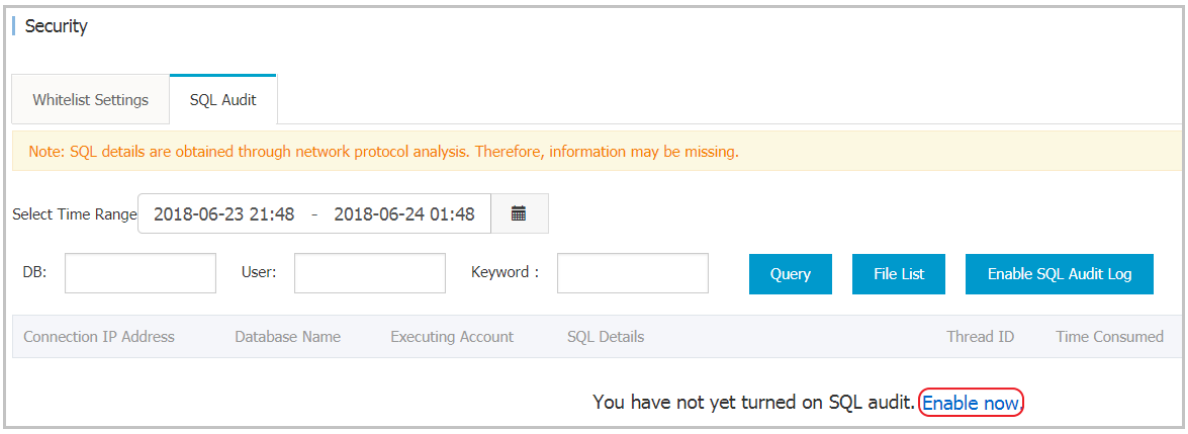

6. In the displayed dialog box, click **Confirm**.

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 them 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, select **Security**.
5. Select the **SQL Audit** tab. Click **Export File** and then click **Confirm**.

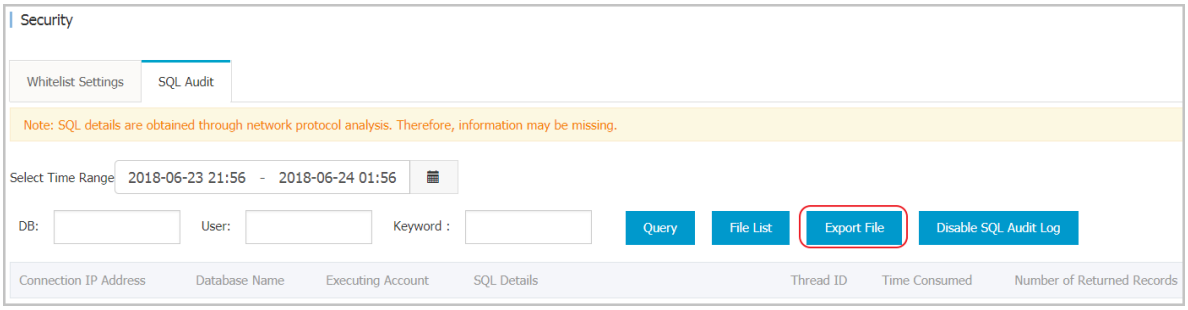

6. Download the SQL audit file and put it in a safe place.
7. Click **Disable SQL Audit Log** and then click **Confirm**.

9.2 Set whitelist

A whitelist is used to allow specified IP addresses and IP segments to access RDS. By default, the RDS 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.

After you set the whitelist, only the following servers can access RDS:

- Servers whose IP addresses are in the whitelist
- ECS instances that are in the security group specified in the whitelist

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*.

We recommend that you periodically check and adjust your whitelists according to your requirements to maintain RDS security. The whitelist only controls access to the RDS instance and does not affect its running. This document describes how to set the whitelist.

Attentions

- The default whitelist group can only be modified or cleared, but cannot be deleted.
- % or 0.0.0.0/0 indicates any IP address is allowed to access the RDS instance. This configurat ion greatly reduces the security of the database and is not recommended.
- When the whitelist is set, the system automatically generates the intranet address for the RDS instance. If you need an Internet address, see *Set intranet and Internet addresses*.
- 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. Select the region where the target instance is located.
3. Click the name of the target instance to go to the **Basic Information** page.
4. Select **Security** in the left-side navigation pane to visit the **Security** page.
5. On the **Whitelist Settings** tab page, find the **default** whitelist group and click **Modify**.

Note:

To add a custom whitelist group to the RDS instance, locate the **default** whitelist group and click **Clear** to delete the IP address 127.0.0.1, and then click **Add a whitelist Group**. The subsequent steps for a custom whitelist are similar to the following steps.

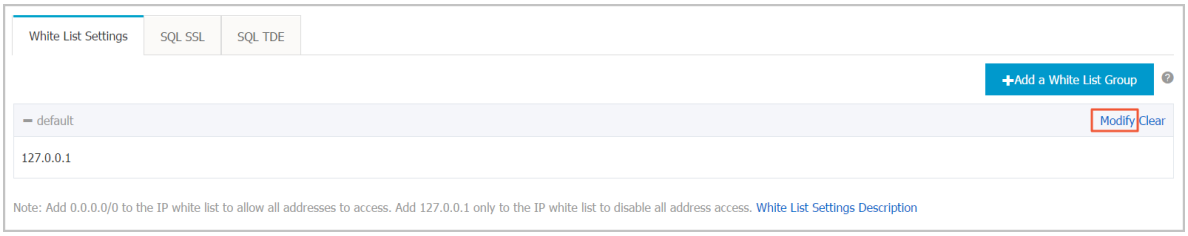

6. On the **Modify Group** page, add the IP addresses or IP segments to access the RDS instance to whitelist field. If you want to add the ECS intranet IP addresses, click **Upload ECS intranet IP Address**, select IP addresses, and click **OK**, as shown in the following figure.

Note:

If you add an IP address or segment to the default group, the default IP address 127.0.0.1 Is automatically deleted.

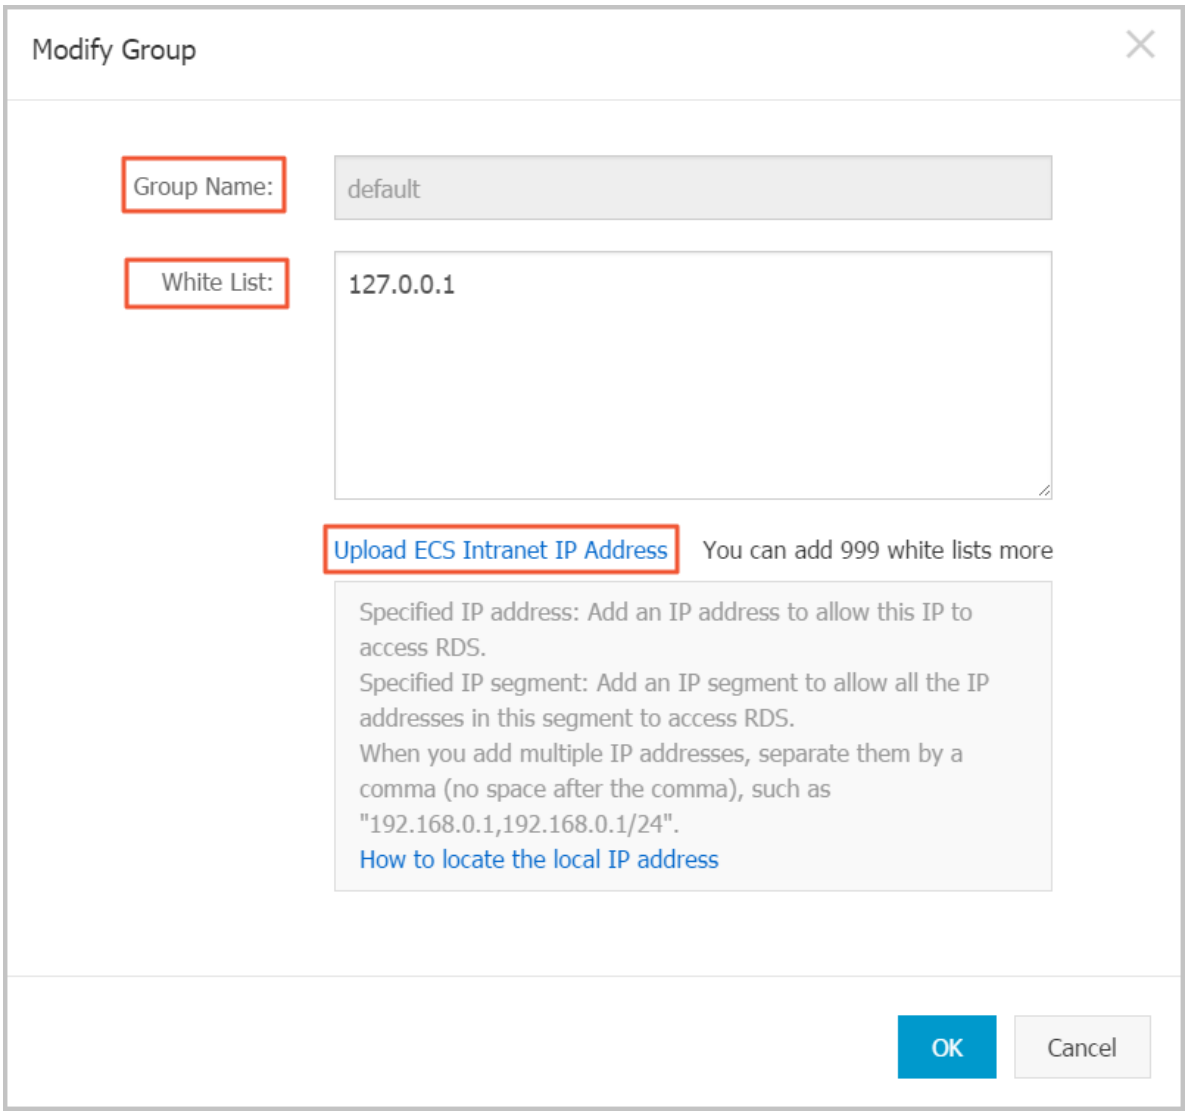

Parameter descriptions:

- **Group Name**: it can contain 2 to 32 characters including lowercase letters, digits, or underscores. The group name must start with a lowercase letter and end with a letter or digit. This name cannot be modified once the whitelist group is successfully created.
- **Whitelist**: enter the custom IP addresses or IP segments that can access the RDS instance.
  - If you enter an IP 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 segments, separate them by commas (,) (do not add blank spaces), such as 192.168.0.1,172.16.213.9.
  - For each whitelist group, up to 1,000 IP addresses or IP segments can be set for MySQL , PostgreSQL, and PPAS instances; and up to 800 can be set for SQL Server instances.

- **Upload ECS intranet IP Address**: by clicking this button, you can select the intranet IP addresses of the ECS instances under the same account with the RDS instance, which is a quick method to add ECS intranet IP addresses.

Precautions for adding ECS security groups

You can configure both the IP whitelist and the ECS security group. Your RDS instance allows access from servers whose IP addresses are in the IP whitelist and ECS instances that are in the security group.

- Currently, only MySQL 5.6 and the Hangzhou, Qingdao, and Hong Kong regions support ECS security groups.
- One RDS instance supports one security group.
- Updates to the ECS security group are automatically applied to the whitelist.

Add an ECS security group

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. Select **Security** in the left-side navigation pane to visit the **Security** page.
5. On the **Whitelist Settings** tab page, click **Add to Security Group**.

> **Note:**
> The security groups marked with "VPC" are in VPCs.

6. Select a security group and click **OK**.

9.3 Set SSL encryption

To increase link security, you can enable SSL encryption and install SSL certificates on the necessary application services. SSL (Secure Sockets Layer) is used on the transport layer to encrypt network connections. It increases the 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. Do this operation 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, select **Security** to go to the **Security** page.
5. Select the **SSL** tab.
6. Click the button next to **Disabled**, as shown in the following figure.

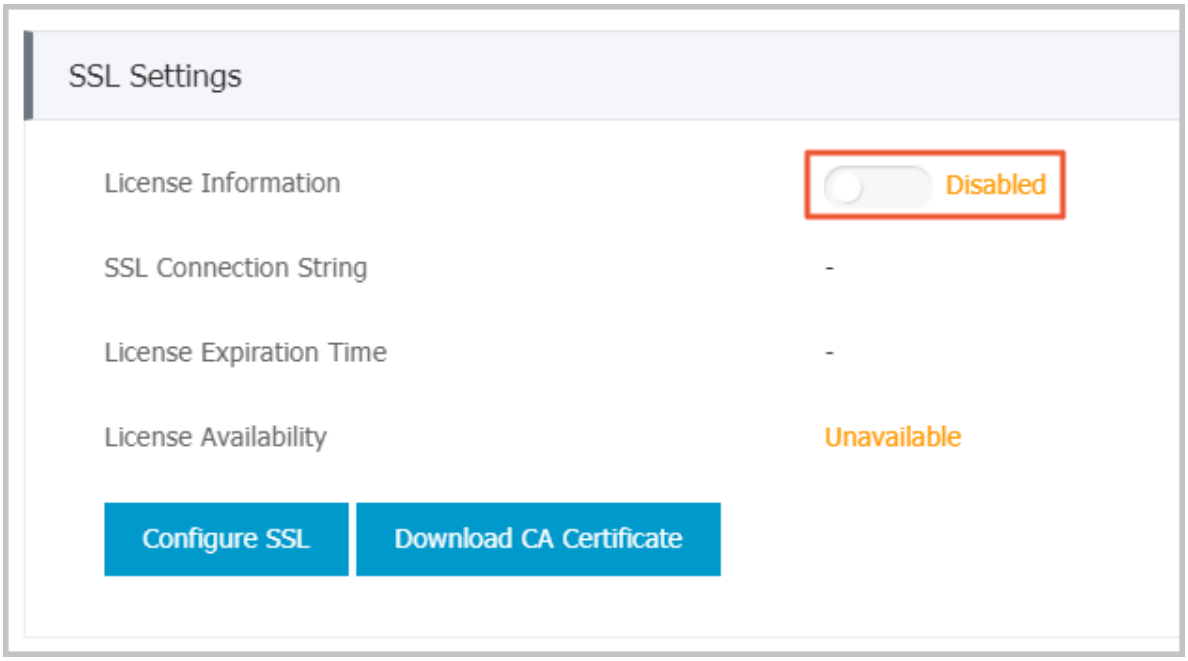

7. In the **SSL Setting** dialog box, select the link for which to activate SSL encryption and click **OK** to activate SSL encryption, as shown in the following figure.

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

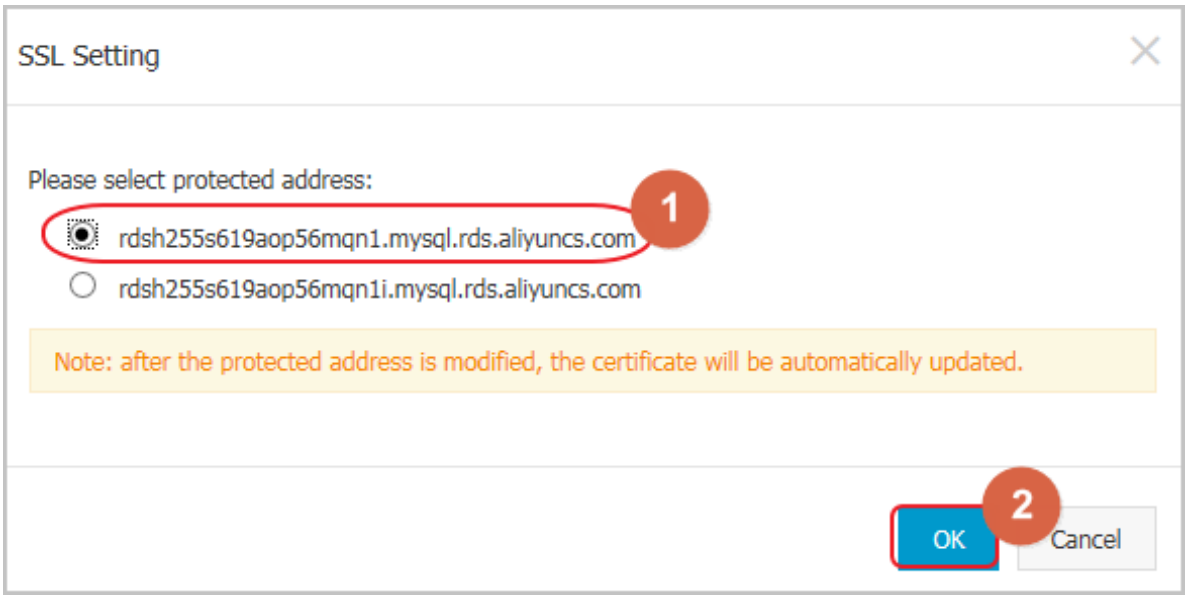

8. Click **Download CA Certificate** to download the SSL certificate, as shown in the following figure.

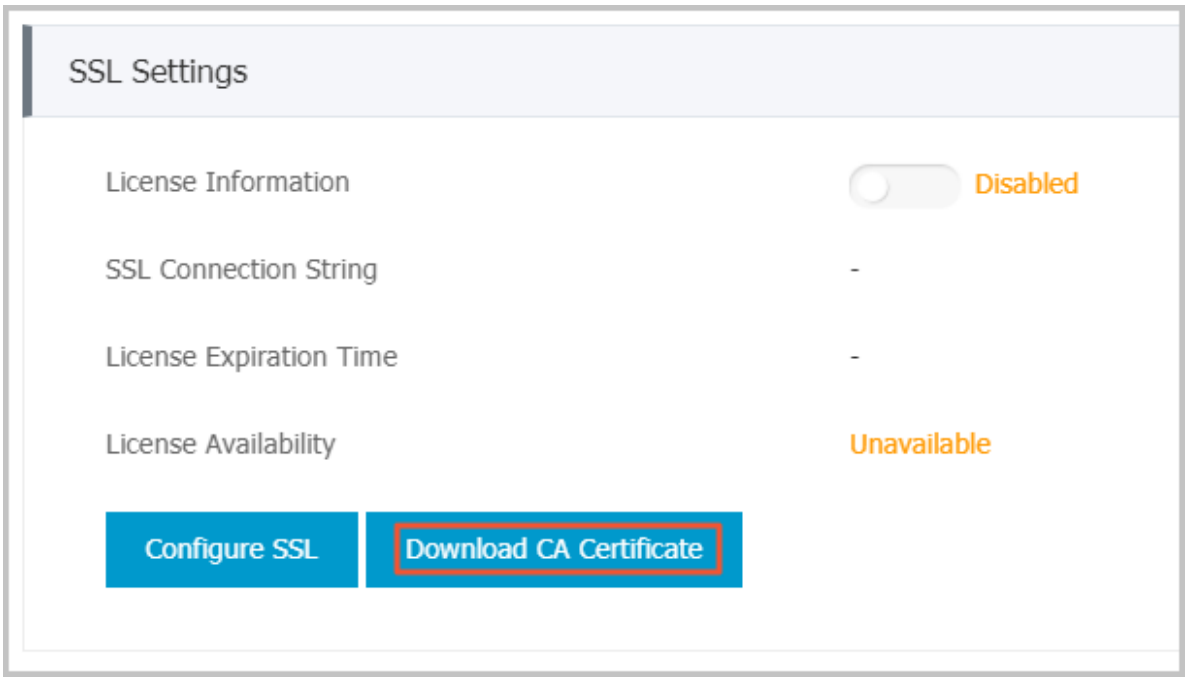

The downloaded SSL certificate is a package including the following files:

- p7b file: Used to import the CA certificate on Windows OS
- PEM file: Used to import the CA certificate on other systems or for other applications
- JKS file: 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, you need to 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.disabledAlgorithms=SSLv3, RC4, DH keySize < 224
jdk.certpath.disabledAlgorithms=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.SSLHandshakeException: DHPublicKey does not comply
to algorithm
            constraints
```

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.

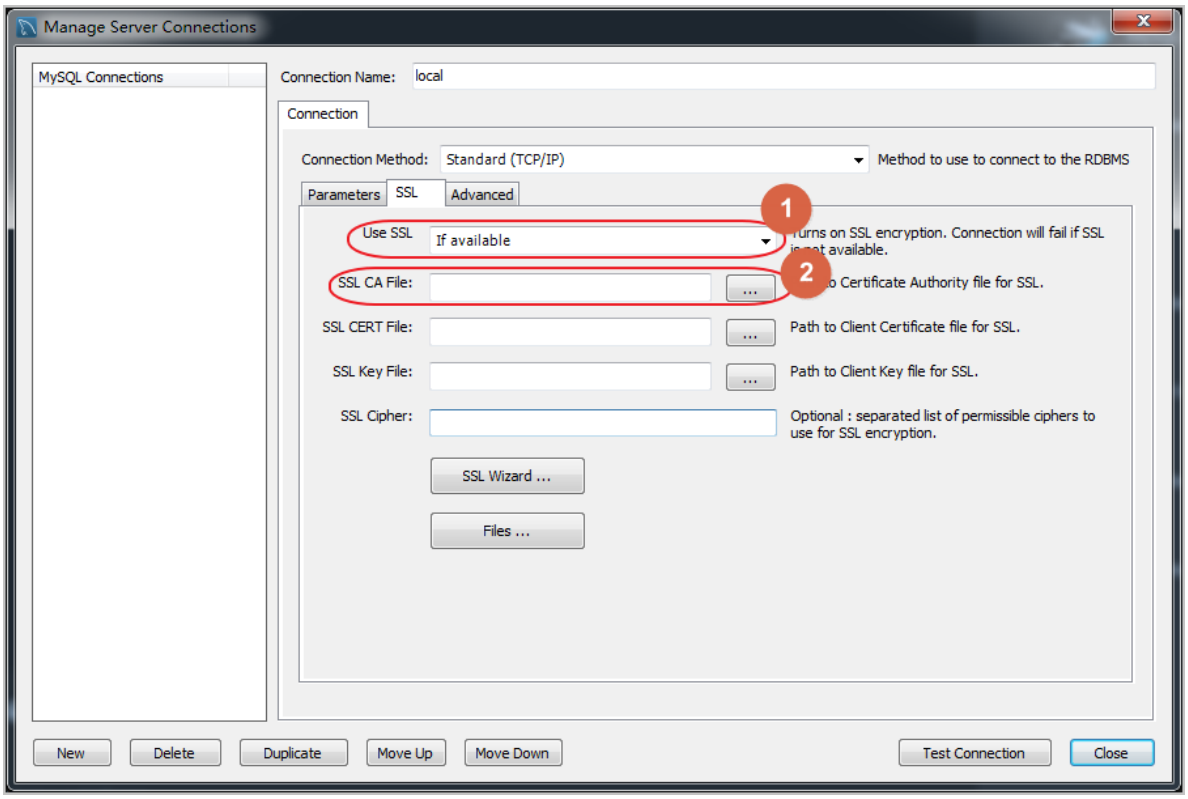

9.4 Set Transparent Data Encryption

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

Note: Currently TDE is only applicable to the database of SQL Server 2008 R2 and MySQL 5.6. To view or modify TDE settings, you need to log in 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 reading from the disk into the memory. TDE does not increase the size of data files. Developers does 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 needed for encryption. After activating TDE, if the user wants to restore the data to the local device, he must use RDS to decrypt the data first.
- After activating TDE, CPU usage significantly increases.

Prerequisite

Key Management Service (KMS) is activated.

Procedure

1. Log on to the *RDS console* and select the target instance.
2. Select **Data Security** in the left-side navigation pane. Then, on the **Data Security** page, select the **TDE** tab.
3. Click **Not Activated**, as shown in the following figure.

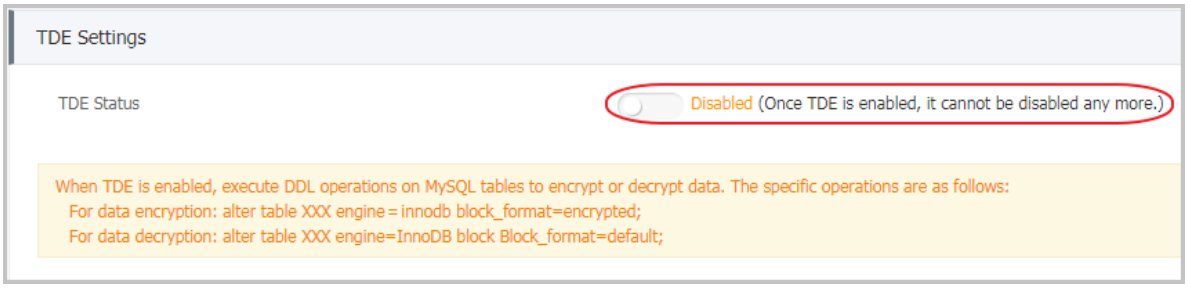

4. Click **OK** to activate TDE.

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

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

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

Subsequent operations

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

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

10 Monitoring and Alarms

10.1 Set monitoring frequency

Background information

The RDS console provides abundant performance metrics for users to conveniently view and know the running status of the 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.

Before May 15, 2018, two monitoring frequencies were available.

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

On May 15, 2018, the second-level monitoring frequency was introduced.

However, the minute-level monitoring frequencies were not enough for certain users and maintenance personnel. Therefore, since May 15, 2018, RDS has introduced the second-level monitoring frequency. This facilitates problem locating and improves customer satisfaction.

- **once per 5 seconds (monitoring period: 7 days) The monitoring frequency beyond 7 days is once per minute.**
- 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), China (Shenzhen)

- The instance is an RDS for MySQL instance.
- The instance storage type is local SSDs.
- The instance memory is 8 GB or more.

- All engines (MySQL, SQL Server, PG, PPAS) and database versions support these 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. Select **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 article.
5. Select the **Monitoring** tab page.
6. Click **Set Monitoring Frequency**.
7. Select the monitoring frequency in the **Set Monitoring Frequency** dialog box and click **OK**.

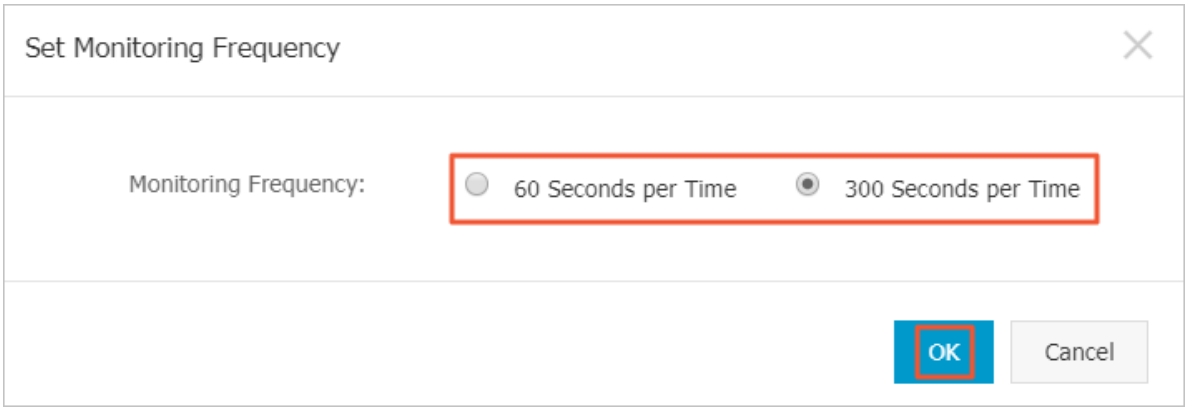

8. If a **Confirm** dialog box is displayed, click **OK**.
9. On the **Monitoring** page, you can also do the following:

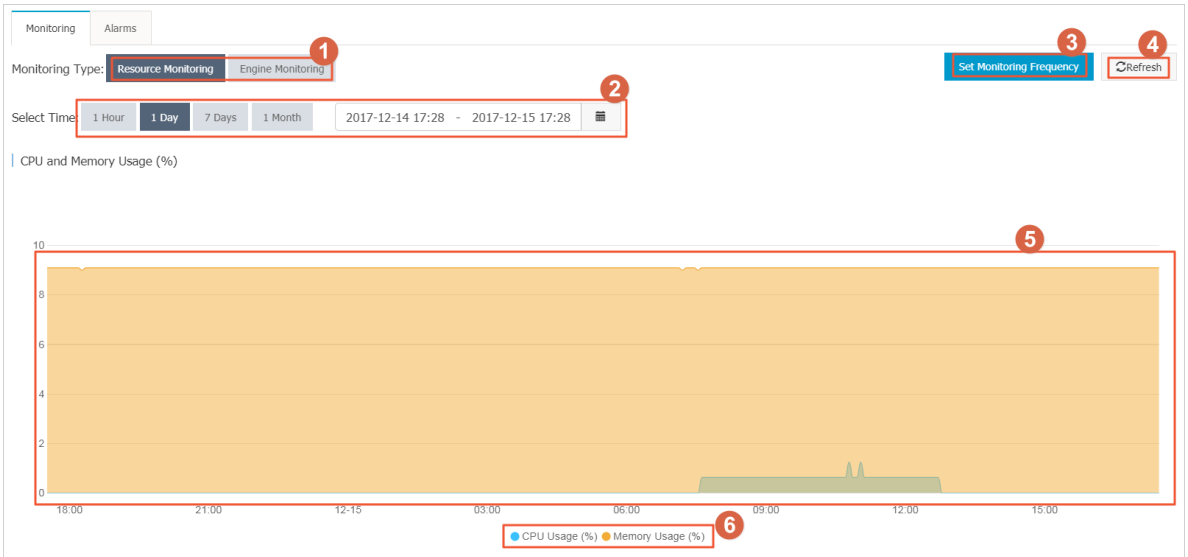

Interface description:

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

List of monitoring items

RDS for MySQL

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: MByte
IOPS	Number of I/O request times of the 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	Usage of CPU and memory of the instance (not including memory used by the operating system).
Network Traffic	Incoming/outgoing traffic of an instance per second. Unit: KByte
QPS/TPS	The number of SQL statements run 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 reads and writes per second. Unit: KByte
The number of InnoDB reads and writes per second.	The number of Read and Write times per second of InnoDB.
InnoDB log	The number of InnoDB physical writes to the log file, log write requests, and FSYNC writes to the log file.
Temporary tables	The number of temporary tables created automatically on the hard disk when the database runs the SQL statement.
MyISAM Key Buffer	Average per-second Key Buffer read hit rate, write hit rate and usage of MyISAM.
MyISAM read and write times	Times of MyISAM read and write from/to the buffer pool and from/to the hard disk per second.
COMDML	The number of statements run for the database per second. The statements include: • Insert • Delete • Insert_Select • Replace • Replace_Select • Select • Update
ROWDML	The number of operations performed on InnoDB, including: • The number of physical writes to the log file per second • The number of rows read in InnoDB tables per second • The number of rows updated, deleted, and inserted in InnoDB tables per second

RDS for SQL Server

Monitoring item	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: MByte
IOPS	I/O request times of the 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 the operating system) of the instance.
Network Traffic	Incoming/outgoing traffic of an instance per second. Unit: KByte
TPS	The number of transactions processed per second.
QPS	The number of SQL statements run per second.
Cache hit rate	Read hit rate of the buffer pool.
Average full table scans per second	Average number of full table scans per second.
SQL compilations per second	The number of compiled SQL statements per second.
Page writes of the checking point per second	The number of page writes of the checking point in the instance per second.
logons per second	The number of logons per second.
Lock timeouts per second	The number of lock timeouts per second.
Deadlocks per second	The number of deadlocks in the instance per second.
Lock waits per second	The number of lock waits per second.

RDS for PostgreSQL

Monitoring item	Description
Disk Space	Usage of the instance disk space. Unit: MByte
IOPS	The number of I/O requests of the data disk and log disk in the instance per second. Unit: time/second

RDS for PPAS

Monitoring item	Description
Disk Space	Usage of the instance disk space. Unit: MByte
IOPS	The number of I/O requests of the data disk and log disk in the instance per second. Unit: time/second

10.2 Set monitoring rules

The RDS instance offers the instance monitoring function, and sends messages to users after detecting an exception in the instance. Besides, when the instance is locked due to the insufficient disk space, the system sends a message to notify users.

Background information

Alibaba CloudMonitor offers monitoring and alarm rules service. CloudMonitor monitors the metrics and you can use the alarm rules service as well. This service helps you to set the alarm rules for the metrics. You must add the alarm contact while forming the contact group. The alarm contact and the contact group is notified immediately when an alarm is triggered in the event of exception. You can create an alarm contact group corresponding to the 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. Select **Monitoring and Alarms** in the left-side navigation pane.
5. Select 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 on 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, select the **Alarm Contact Group** tab.
11. Click **Create Alarm Contact Group**.
12. Fill in **`Group Name`** and **`Description`**, select a contact from the **Existing Contacts**, and 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 member in the contact group.

13. After creating the alarm contact group, select **Cloud Service Monitoring** > **> ApsaraDB for RDS** in 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 Log management

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

Use the console to manage logs

You can use the console to manage logs of MySQL 5.5/5.6, SQL Server 2008 R2, PostgreSQL, and PPAS instances. The actual interface may vary depending on the engine type and version.

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. Select **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 statement that failed to be executed in the past month.
Slow SQL Log Details	• Records the SQL statements that lasted for over 1 second (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 1 second (You can modify this time threshold by modifying the long_query_time parameter in Parameters) in the past month.

Use SQL commands to manage logs

To query error logs of SQL Server 2012 and later versions, use this procedure:

```
sp_rds_read_error_logs
```

The usage method is the same as using this stored procedure:

```
sp_readerrorlog
```

Examples are as follows:

Example 1:

```
EXEC sp_rds_read_error_logs
```

Example 2:

```
EXEC sp_rds_read_error_logs 0,1 ,'error'
```

12 Linked server of SQL Server instances

This article is applicable only to high-availability instances of RDS SQL Server 2012 and later versions.

Currently, linked server creation has the following constraints:

- You cannot create a linked server on the RDS console.
- Creating a linked server with a series of stored procedures is complicated.
- You cannot create a linked server by using DNS and the corresponding IP address.

Despite the constraints, this article provides a simple method for creating a linked server.

```
DECLARE
        @linked_server_name sysname = N'my_link_server',
        @data_source sysname = N'***********', --style: 10.1.10.1,1433
        @user_name sysname = N'****' ,
        @password nvarchar(128) = N'**********',
        @link_server_options xml
        = N'
            <rds_linked_server>
                <config option="data access">true</config>
                <config option="rpc">true</config>
                <config option="rpc out">true</config>
            </rds_linked_server>

        EXEC sp_rds_add_linked_server
            @linked_server_name,
            @data_source,
            @user_name,
            @password,
            @link_server_options
```

The following message “create successfully” appears after the linked server is successfully created.

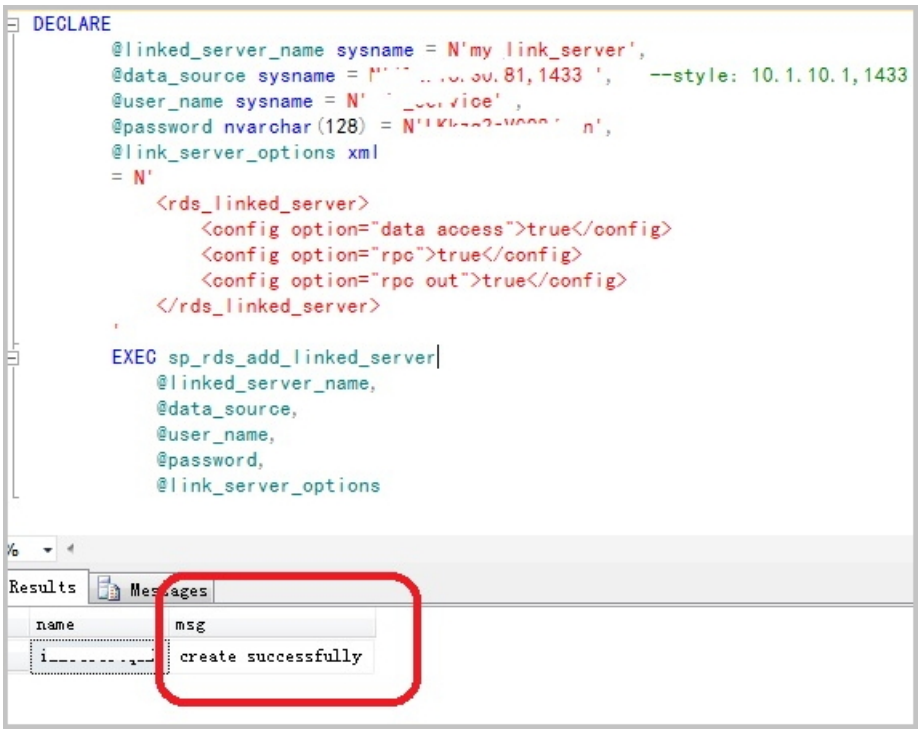

Click the **Messages** tab shown in the preceding figure, and the following information is displayed.

```
The linked server 'my_link_server' has set option 'data access' to '
true'.
The linked server 'my_link_server' has set option 'rpc’ to 'true'.
The linked server 'my_link_server' has set option 'rpc out' to 'true'.
create link server 'my_link_server' successfully.
```

13 Backup and recovery

13.1 Recover MySQL data

13.1.1 Recover data from a clone instance to a master instance

The data recovery function minimizes the damage caused by database misoperations. We recommend that you recover data to the master instance through a clone instance.

Currently, the following RDS instances support clone instances.

- MySQL 5.5, 5.6, 5.7 master instances (except MySQL 5.7 Basic Edition)
- SQL Server 2016 High-Availability Edition (including Standard and Enterprise Editions)
- SQL Server 2012 High-Availability Edition (including Standard and Enterprise Editions)

To recover data for other RDS instances, see *Recover data from a temporary instance to a master 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 to go to the **Basic Information** page.
4. Recover data to a clone instance. For more information, see *Create a clone instance*.
5. After the clone instance is created, go back to the **Basic Information** page of the master instance.
6. Click **Create Data Migration Task** at the top of the page.
7. Select **Data migration** in the left-side navigation pane.
8. Click **Create migration task** at the upper right corner.
9. Enter the task name, source database information, and target database information.

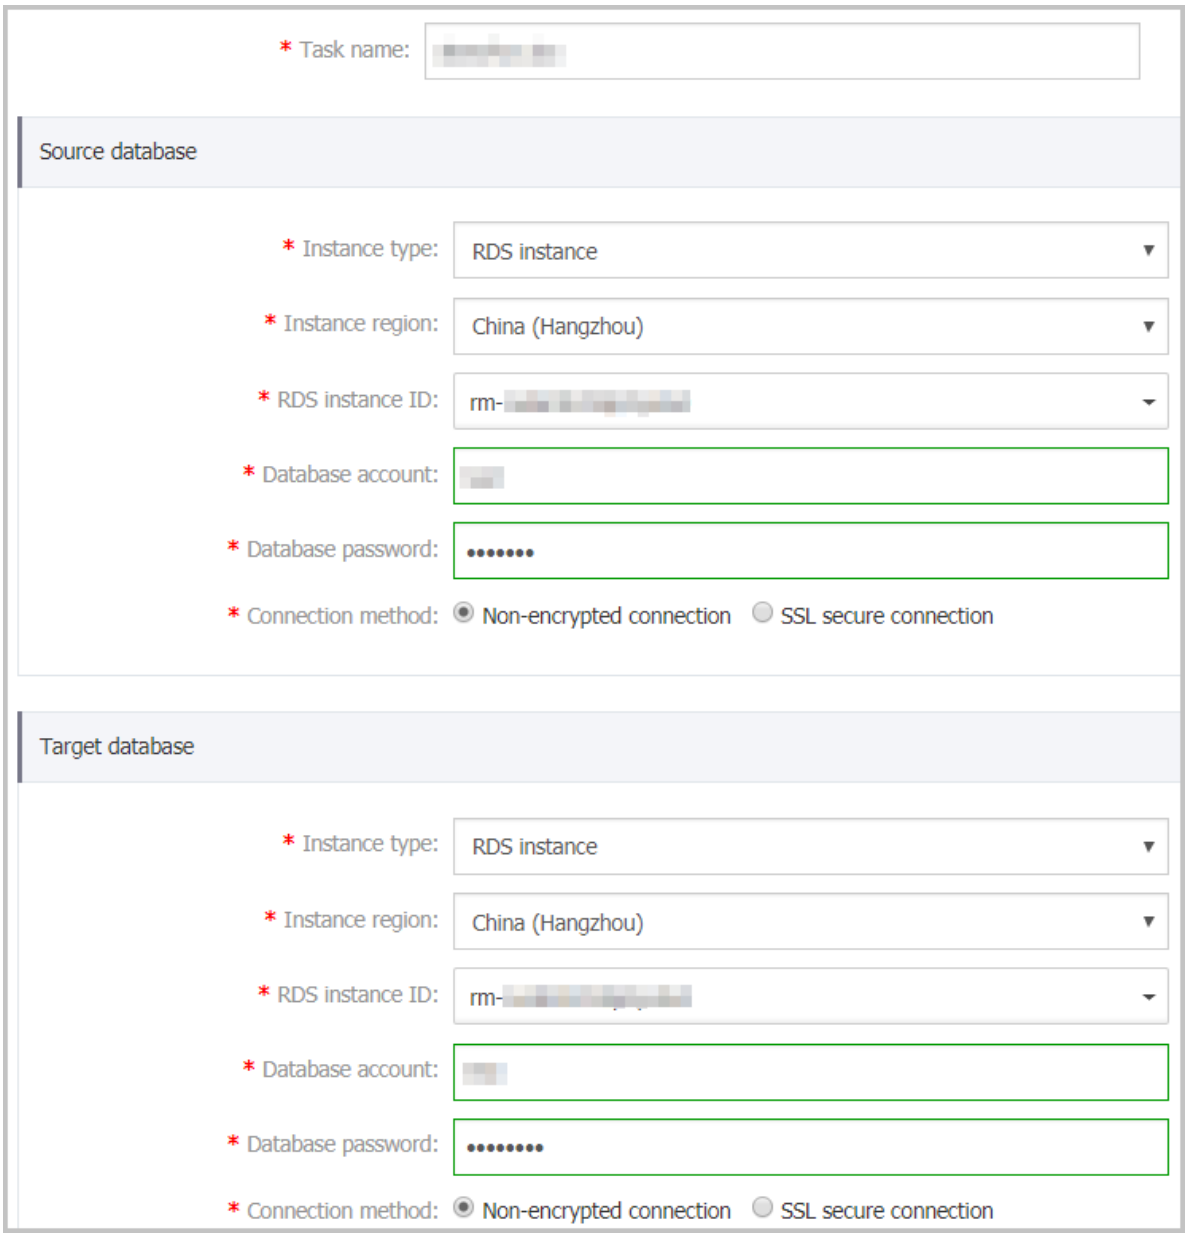

Parameter descriptions:

- By default, DTS automatically generates a name for each task. You can edit the name to indicate the specific services for easy identification of the task.
- Source database information:
  - Instance type: Specifies the instance type of a database. Select **RDS Instance**.
  - Instance region: The region where the master instance is located.
  - RDS instance ID: Click the drop-down list and select the clone instance ID.
  - Database account: It is consistent with the account name of the master instance. Make sure that this account has the read/write privilege to all the data to be migrated.

- Database password: It is consistent with the password of the master instance account.

- Target database information
  - Instance type: **RDS Instance** by default.
  - Instance region: The region where the master instance is located.
  - RDS instance ID: Click the drop-down list and select the master instance.
  - Database account: The master instance account name. Make sure that this account has the read/write privilege to all the data to be migrated.
  - Database password: The password of the master instance account.

10. Click **Authorize whitelist and enter into next step**.
11. Select the migration type, choose objects from the **Migration objects** column, and click **>** to add the objects to the **Selected** column.

    To modify the migration object name in the target database, you can click **Edit** on the right side of the **Selected** list to modify the name.
12. Click **Pre-check and start**.
13. If the system displays the pre-check failure result, click **!** next to the check item with **Check Result**as **Failed** to check the detailed failure information, and perform troubleshooting accordingly.
14. After troubleshooting, select the current migration task on the **Migration task list** page and click **Start**.
15. After pre-check is passed, click **OK** to automatically run the migration task.

13.1.2 Recover data directly to the master instance

Procedure

You can recover data directly to the master instance, and the specified backup data overwrites the data of the master instance, but the data generated after creation of the specified backup data is lost. We recommend that you create a temporary instance for data recovery and migration to guarantee higher security.

Note:

- This article is applicable only to MySQL instances. To recover data directly to a SQL server instance, see *Recover data directly to the master instance*

- If a read-only instance exists, the specified backup data cannot directly overwrite the original data of the master instance. You can only recover the data through a clone instance. For more information, see *Recover from a clone instance to the master instance*.

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. Select **Backup and Recovery** in the left-side menu.
5. Select the **Backup List** tab.
6. Select the time range and click **Query**.
7. Find the target backup and click **Restore** in the **`Action`** column.
8. In the displayed dialog box, select **Coverage Restoration** and click **OK**.
9. Click **OK** again.

13.2 Recover SQL Server/PPAS/PostgreSQL data

13.2.1 Recover data to the master instance through a temporary instance

Note:

This article is not applicable to MySQL instances. To recover data for a MySQL instance, see *Recover data from a clone instance to a master instance* (recommended) or *Recover data directly to the master instance*.

The data recovery function minimizes the damage caused by database misoperations. We recommend that you recover data to the master instance through a temporary instance. That is to say, 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.

Note:

- 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 any 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. Select **Backup and Recovery** in the left-side navigation pane.
4. Select 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. Select **Data migration** in the left-side navigation pane.
4. Click **Create migration task**.
5. Set the parameters.
   - Task name: A default task name is generated. You can modify it so that you will 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: Same as the account name of the master instance. Make sure that this account has read/write permissions on the data to be migrated.
     - Database password: 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 the temporary instance.

- Database account: Enter the account name of the master instance. Make sure that this account has read/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 types.
8. In the left pane, select the objects to be migrated and click **>** to add them to the right. 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 start the pre-check again.
11. After the pre-check is passed, click **OK** to start the migration task.

13.2.2 Recover data directly to an instance

You can recover data directly to an instance, and the specified backup data overwrites the data of the instance, and the data generated after creation of the specified backup data is lost. We recommend that you create a temporary instance for data recovery and migration to guarantee higher security.

> **Note:**
> This method is only applicable to the database of SQL Server 2008 R2.

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. Select **Backup and Recovery** in the left-side menu.
5. Select the **Backup List** tab.
6. Select the time range for recovery and click **Query**.
7. Select the target backup file and click **Coverage Restoration**, as shown in the following figure.

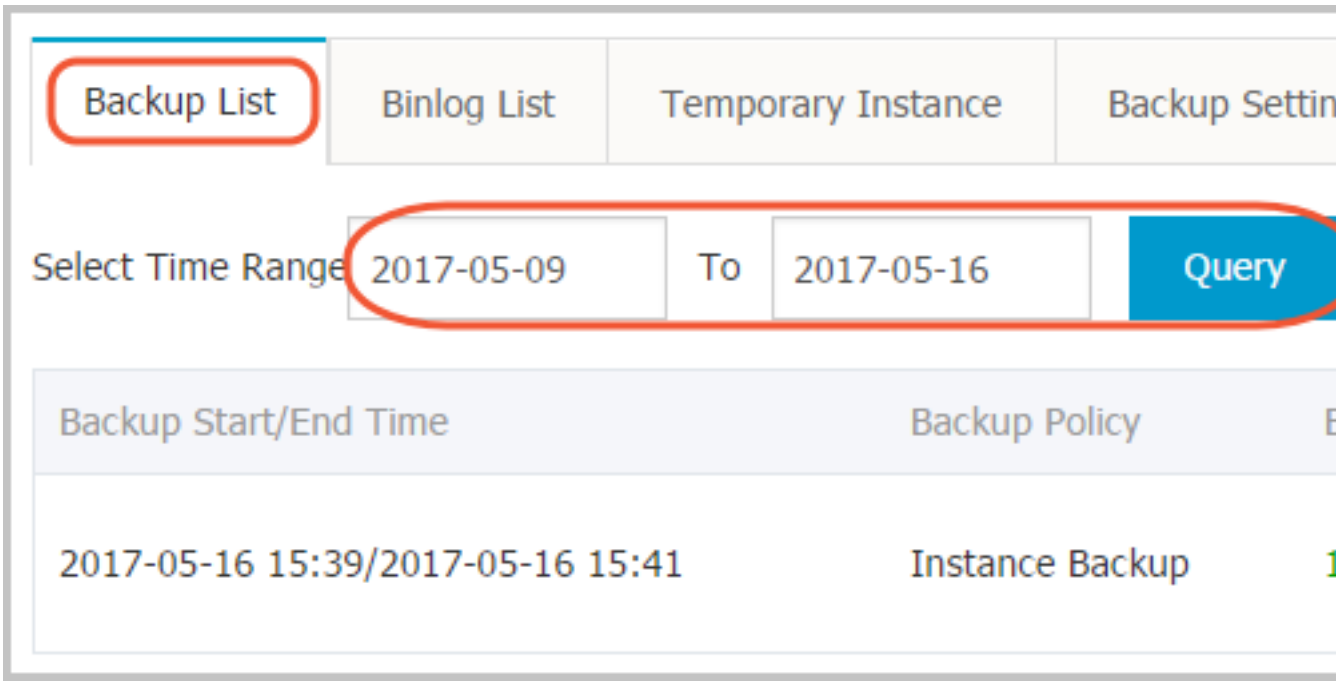

8. Click **Confirm** in the dialog box to recover data to the master instance.

13.3 Back up RDS data

You can configure a backup policy to adjust the cycles of RDS data backup and log backup and 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 design a backup cycle appropriately to cater the service requirements based on the available backup space. For information on 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 by 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.	• 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

Database type	Data backup	Log 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, not logical backup. — Backup files are retained for at most 7 days for free.	files to the OSS. This does not affect the data recovery function and stops the Binlog files from occupying instance disk space.
SQL Server	• Supports full physical backup and incremental physical backup. • Automatic backup uses the cycle: Full Backup-Incremental Backup-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, and so on. If a full backup is manually performed at any time in the backup cycle, the next two backups are incremental backups. • The 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 the transaction log.	Included in data backup; individual transaction logs are not provided for download.
PostgreSQL	Supports full physical backup.	After being generated, write-ahead logs (WALs) (16 MB per log) are compressed and uploaded immediatel y. Local files are deleted within 24 hours.

Database type	Data backup	Log backup
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.

Automatic backup (Backup policy setting)

After you configure a backup policy, RDS automatically backs up databases based on the policy.

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. Select **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:

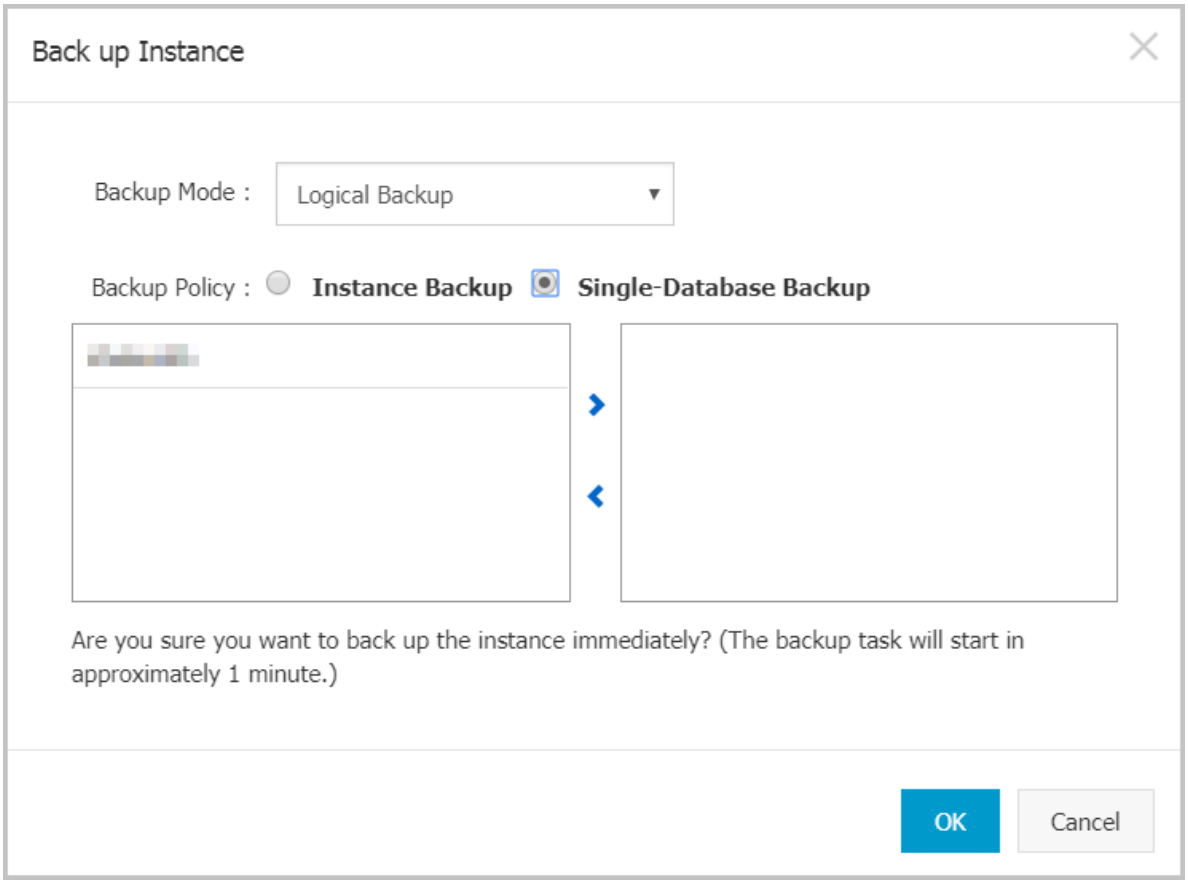

Parameters	Description
Data Retention Period (days)	• Specifies the time period for which backup files are retained. The default value is 7 days. The value range is 7 to 730 days. • MySQL 5.7 Basic Edition backup files are retained for free for at most 7 days.
Backup Cycle Frequency	• You can set it to one or multiple days in a week. • SQL Server, PostgreSQL, and PPAS instances are backed up daily by default and this cannot be modified.
Next Backup	This value can be set to any time; units: hours.
Log Backup	**Enable** or **Disable**
Log Retention Period (days)	• Specifies the number of days when the log backup files are retained. The default value is 7 days. • The value range is 7-730 days and it must be less than or equal to the value of the retention days.

Manual backup

Note:

The following procedure describes how to configure the single-database logical backup for MySQL 5.7 Basic Edition.

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**.

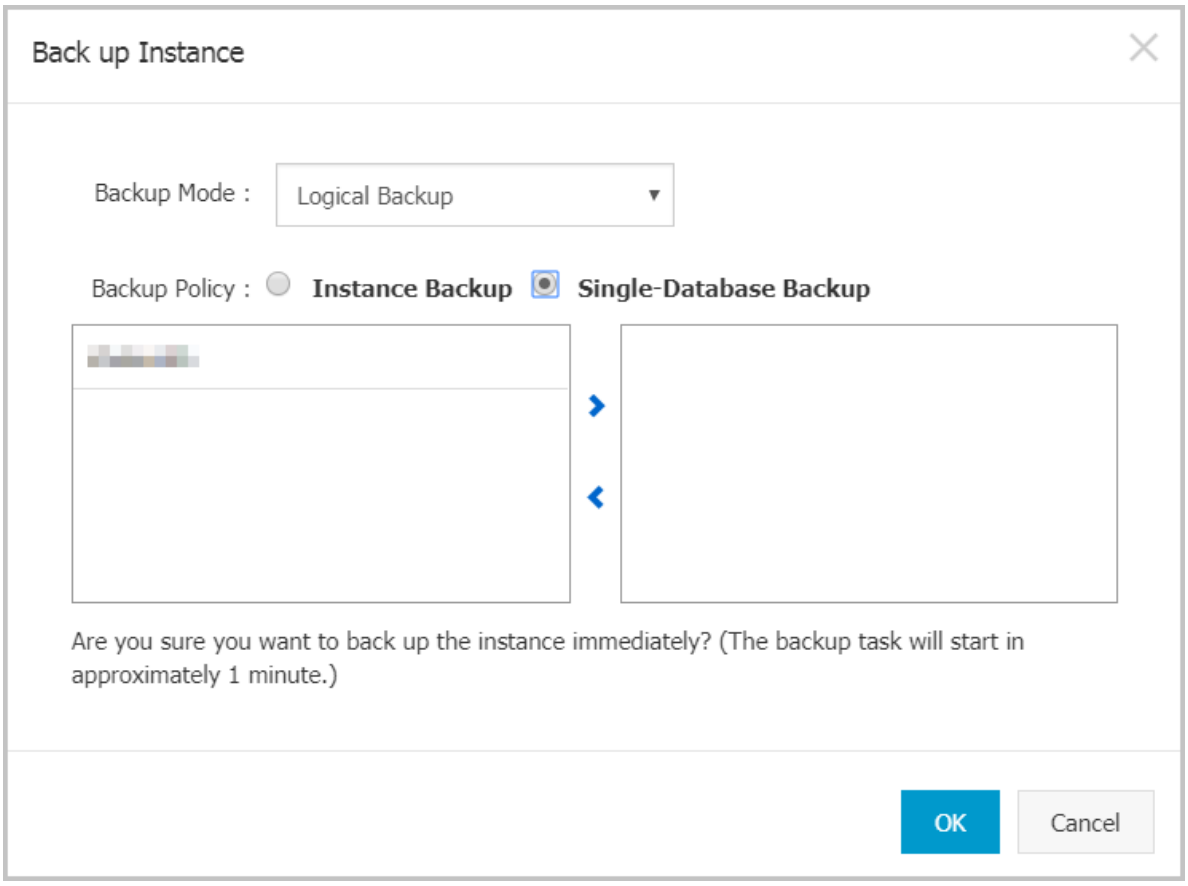

Note:

- 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**.

13.4 View the free quota of the backup space

Backup files of an instance occupy the backup space. Each ApsaraDB for 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 on the billing standard 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 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 ApsaraDB for 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.

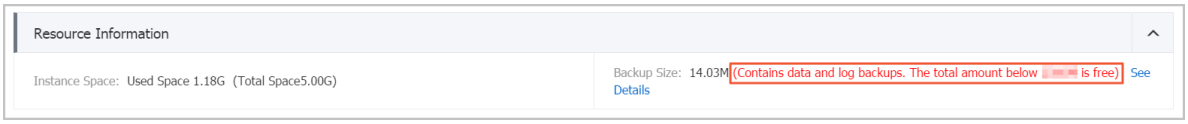

13.5 Download RDS data and log backup

To protect users' rights, RDS allows users to download data backup files 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 menu to enter the **Backup and Recovery** page.
5. Do the following to download a data backup or a log backup:
   - 中国站
     1. 中国站

2. 中国站
3. 中国站

> **Note:**
> 中国站

4. 中国站

中国站	中国站
中国站	中国站
中国站	中国站 中国站
中国站	中国站 中国站

- 中国站

1. 中国站
2. 中国站
3. 中国站

Note:

中国站

- 中国站
- 中国站

4. In the BINLOG File Download confirmation box, select the download method.

Download methods	Note:
Download:	The backup file is downloaded using an Internet address.
Copy intranet address:	中国站 When ECS and RDS are in the same region, you can use an intranet address on ECS to download the backup file at a faster speed and with higher security.
Copy Internet address:	中国站 An Internet address is copied and used to download the backup file by other tools.

13.6 Logical backup and recovery for PPAS

This chapter describes the steps for logical backup and recovery from the RDS instances for PPAS.

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.

   Windows users: *http://yunpan.taobao.com/s/2Y03fmh7PF0* (Access code: VAXVAc).

   Linux users: *http://yunpan.taobao.com/s/1H1T5Kqog8s* (Access code: 561TH4).

2. Grant the permissions of all roles 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 permissions 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>
   ```

FAQs

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

   ```
   ERROR: permission denied for relation product_component_version
   ```

```
LOCK TABLE sys.product_component_version 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 PPAS binary system to export the data. For PPAS downloading methods, see the proceding steps.

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, namely, 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 the Windows platform, 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, for example,

```
pg_dump -Uxxx -h yyy.
```

This parameter is not allowed since a white space is needed next to -U (other parameters also follow this style).

14 Tag management

14.1 Create tags

If you have a large number of instances, you can bind tags to facilitate classification and management. Each tag is composed of a key-value pair, enabling you implement two-level classification for the instances you created.

Constraints

- Up to 10 tags can be bound to a single instance and they must have unique TagKeys. Tags with the same TagKeys are overwritten.
- You cannot bind/unbind more than 5 tags immediately.
- Tag information is independent in different regions.
- After unbinding a tag, if it is not bound to any other instances, it is deleted.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Add tags to the instances. There are two methods:
   - Add tags to a single instance: Locate an instance, and click **More** > **Edit tags** .
   - Add tags to multiple instances: Select multiple instances, and click **Edit tags**.

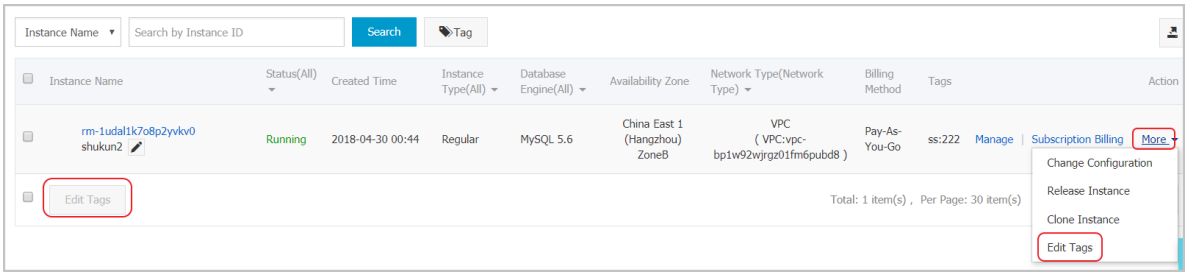

4. If you want to add a new tag, click **Create**. Enter the tag `Key` and `Value`, and click **Confirm**.

> **Note:**
> If you want use the existing tags, click **Available Tags**, select the tag key and tag value, and click **Confirm**.

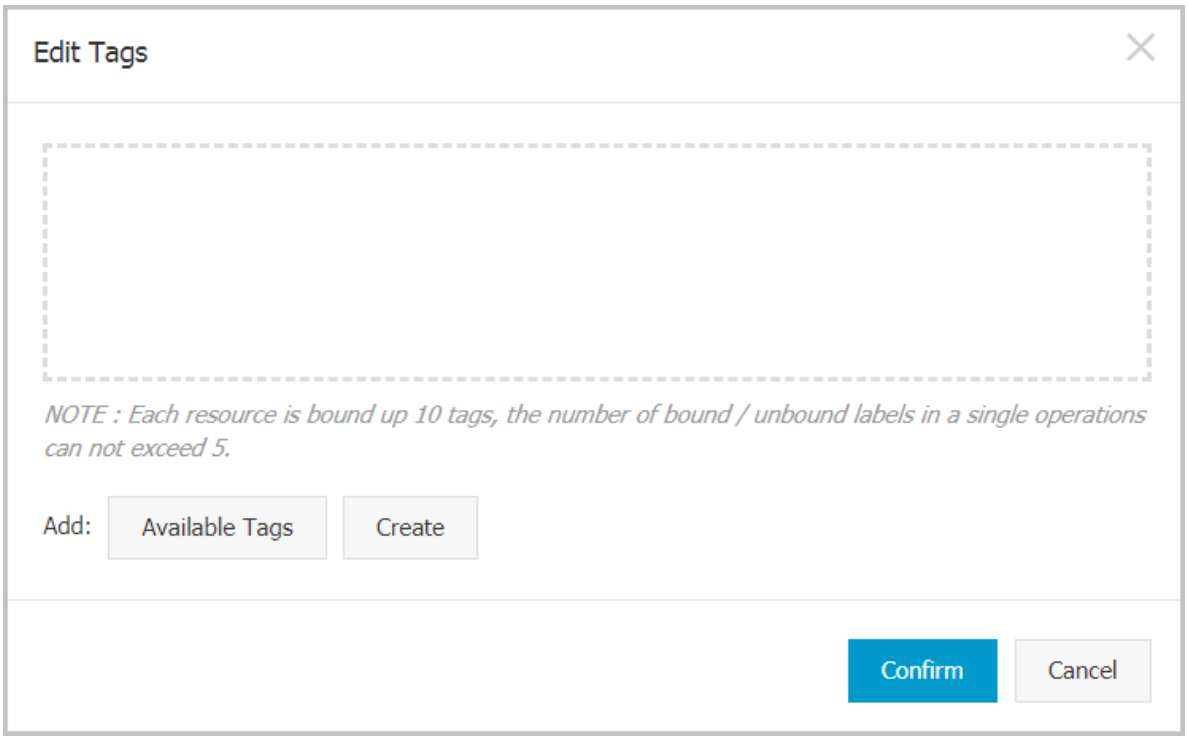

5. When the tags you need have been added to the tag table, click **Confirm**.

14.2 Delete tags

If you have changed your instance or no longer need to use tags, you can delete the tags from the instance.

Constraints

- You cannot bind/unbind more than 5 tags immediately.
- After unbinding a tag, if it is not bound to any other instances, it is deleted.

Procedure

1. Log on to the *RDS console*.
2. Select the region where the target instance is located.
3. Select **More** > **Edit tags** in the **Action** column of the target instance.
4. Click **x** after the tag to delete it, as shown in the following figure.

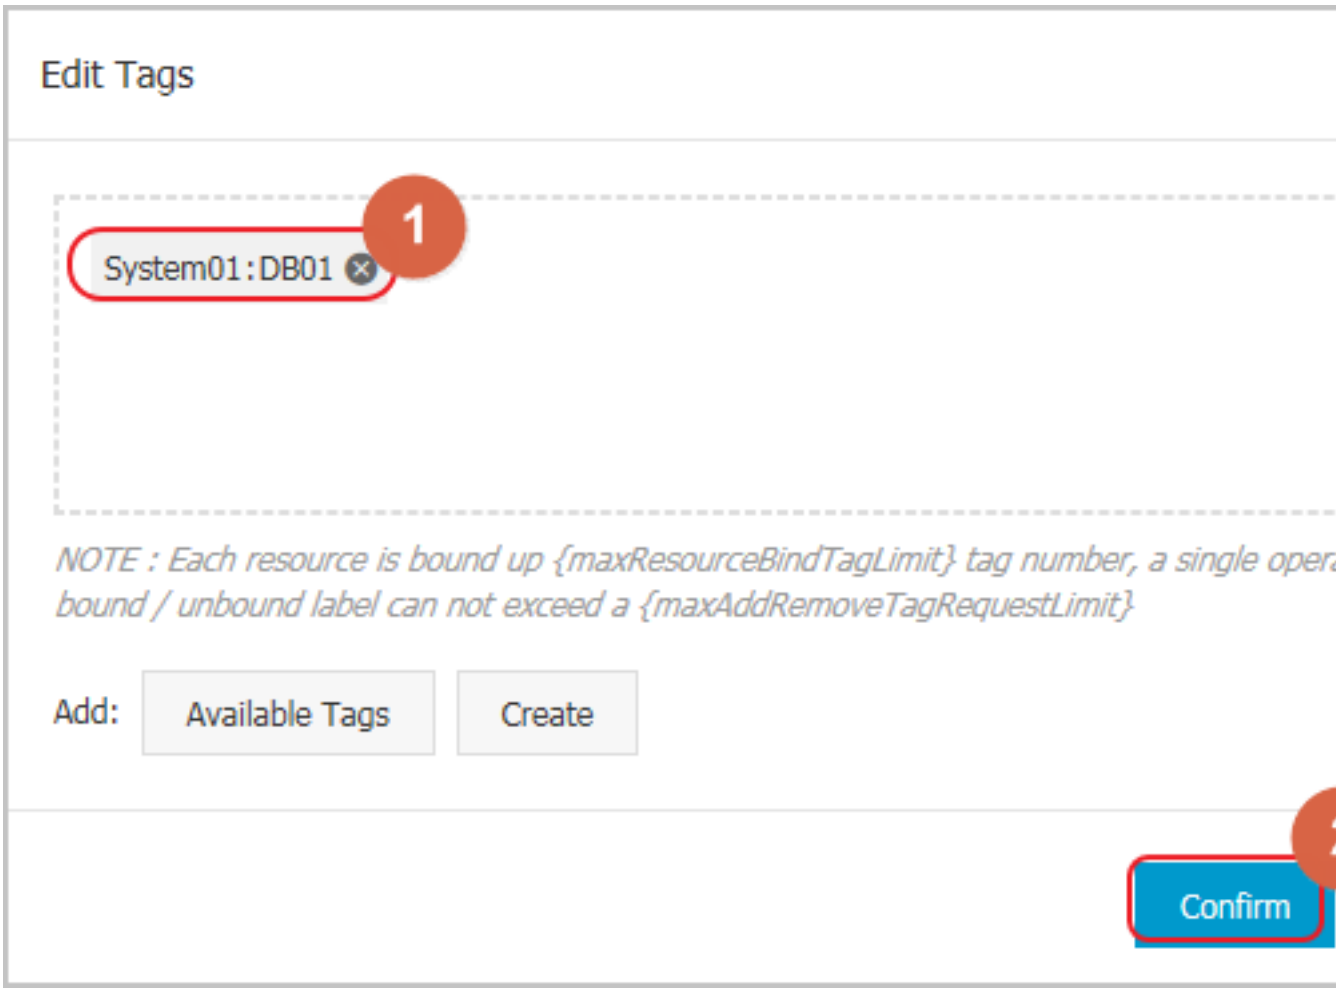

5. Click **Confirm**.

14.3 Filter instances by tag

1. Log on to the *RDS console* to go to the **Instances** page.
2. Click **Tag** next to the **Search** button. Select **Tag Key** and **Tag Value** to filter instances, as shown in the following figure.

> **Note:**
> To cancel a filter, you can remove the selected tag next to the **Tag** button.

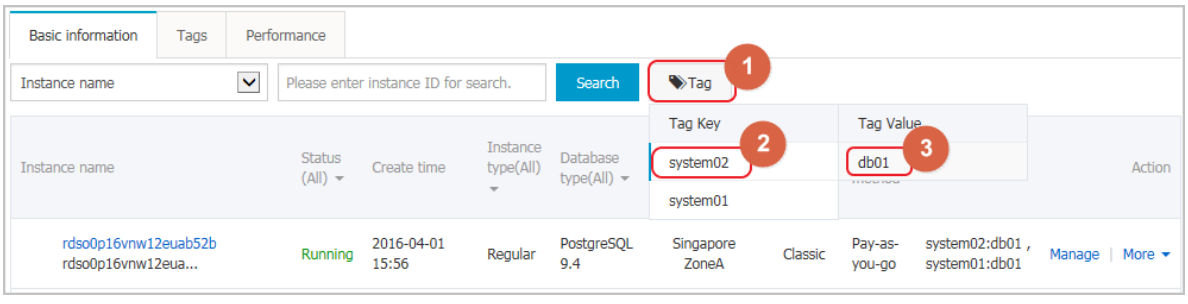

15 Data migration

15.1 中国站目前没有

15.2 Use mysqldump to migrate MySQL data

The mysqldump command 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 the original database to an RDS instance is similar to the procedure for migrating data from one MySQL server to another MySQL server.

Prerequisites

- You have set the whitelist, apply for an Internet address, and create 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 the read/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 hosts, you can use the wildcard `%`.
   - password: Indicates the logon password for the account.

In the following example: Username is *William*; 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 databasename.tablename TO 'username'@'host' WITH
GRANT OPTION;
GRANT REPLICATION SLAVE ON databasename.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*.
- databasename: 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 hosts, you can use the wildcard *%*
- WITH GRANT OPTION: An optional parameter that enables the account to use the GRANT command.

In the following example, the account *William* is granted with 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 exports 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” so as to be compatible 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 article illustrates how to upload files to the following path.

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

6. Log on to ECS and import the 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 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.

15.3 Migrate RDS data to the local database

15.3.1 Migrate RDS for PPAS to local Oracle

Constraints

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-2ze466l5u1k657yyn.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-2ze466l5u1k657yyn.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 information.
 ora=>
```

If you can log on to RDS for PPAS as ora, it means that the connection has been established.

After a successful test, return to the root user.

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

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

   a. Create a file ppas_exp_all_tables_to_csv.sh.

```
 vi ppas_exp_all_tables_to_csv.sh
```

   b. Insert the following text into the ppas_exp_all_tables_to_csv.sh script.

```
# ppas_exp_all_tables_to_csv.sh <hostname> <port> <username> <
database>
 # Author: Xiao Shaocong (Scott Siu)
 # Email: shaocong.xsc@alibaba-inc.com
 TMP_PATH="/tmp/ppas_tables_$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_tables_$1_$2_$3_$4.sql
 psql -h $1 -p $2 -U $3 $4 -f /tmp/ppas_tables_$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_all_tables_to_csv.sh script.

```
 [root@oraclexe ~]# chmod 0755 ppas_exp_all_tables_to_csv.sh
```

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

```
 [root@oraclexe ~]# ./ppas_exp_all_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_tables_rm-2ze466l5u1k657yyn.ppas.
rds.aliyuncs.com_3433_myadmin_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 Oracle.

   - 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 that a directory cannot be created as follows.

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

```
mkdir: Cannot create directory: "/tmp/ppas_tables_rm-2ze466l5u1k657yyn
.ppas.rds.aliyuncs.com_3433_myadmin_ora": file already exists
```

Handling process

Delete the existing directory.

```
[root@oraclexe ~]# rm -rf /tmp/ppas_tables_rm-2ze466l5u1k657yyn.ppas.
rds.aliyuncs.com_3433_myadmin_ora
```

15.3.2 Migrate RDS for MySQL data to the local MySQL database

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

Export using 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, you need first of all recover data to Linux and then migrate the data to Windows.

Prerequisites

RDS adopts the open source software Percona XtraBackup 2.0.6 to perform full physical backup on the MySQL database. You must download the software for data recovery. Visit the official website (*http://www.percona.com/*) of Percona XtraBackup and download the version compatible with your operating system. For example: Download the RHEL6/x86_64 version and run the rpm command to install it.

```
sudo rpm -ivh percona-xtrabackup-2.0.6-521.rhel6.x86_64.rpm
```

Procedure

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

1. Download the RDS **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 download 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 the 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 db0dz1rv11f44yg2
-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_info
-rw-r--r-- 1 root root        77 Aug 19 18:15 xtrabackup_checkpoints
-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_info
```

5. Recover the data file.

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

Data is successfully recovered when the system displays `innobackupex: completed OK!`

6. Modify the configuration file. In the *backup-my.cnf* file, comment out *innodb_fast_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 innobackupex-1.5.1.
# The MySQL Server
[mysqld]
innodb_data_file_path=ibdata1:200M:autoextend
innodb_log_files_in_group=2
innodb_log_file_size=524288000
#innodb_fast_checksum=0
#innodb_page_size=16364
#innodb_log_block_size=512
datadir=/home/mysql/
```

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

```
rm -rf mysql
```

```
mysql_install_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_safe --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           |
+--------------------+
| information_schema |
| db0dz1rv11f44yg2   |
| mysql              |
| performance_schema |
| test               |
+--------------------+
```

Export using a logical backup file

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

Procedure

1. Download the RDS **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 download 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.

15.3.3 Migrate RDS for SQL Server data to the local SQL Server database

RDS for SQL Server supports the migration of cloud data to the local database by 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 download address.

2. After download, 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: Specifies full backup, such as **rdsumu2myfzbeai1_datafull_201402250050_2250050.bak**
   - datadiff: Specifies incremental backup, such as **rdsumu2myfzbeai1_datadiff_201402260050_2260050.bak**

- log: Specifies log backup, such as **rdsumu2myfzbeai1_log_201402260050_2260050.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 to the full backup file: d:\backup\rdsumu2myfzbeai1_datafull_201402250050_2250050.bak
   - Path to the incremental backup file: d:\backup\rdsumu2myfzbeai1_datadiff_201402260050_2260050.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 filelistonly from disk=&apos;d:\backup\rdsumu2myfzbeai1_datafull_201402250050_2250050.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**:

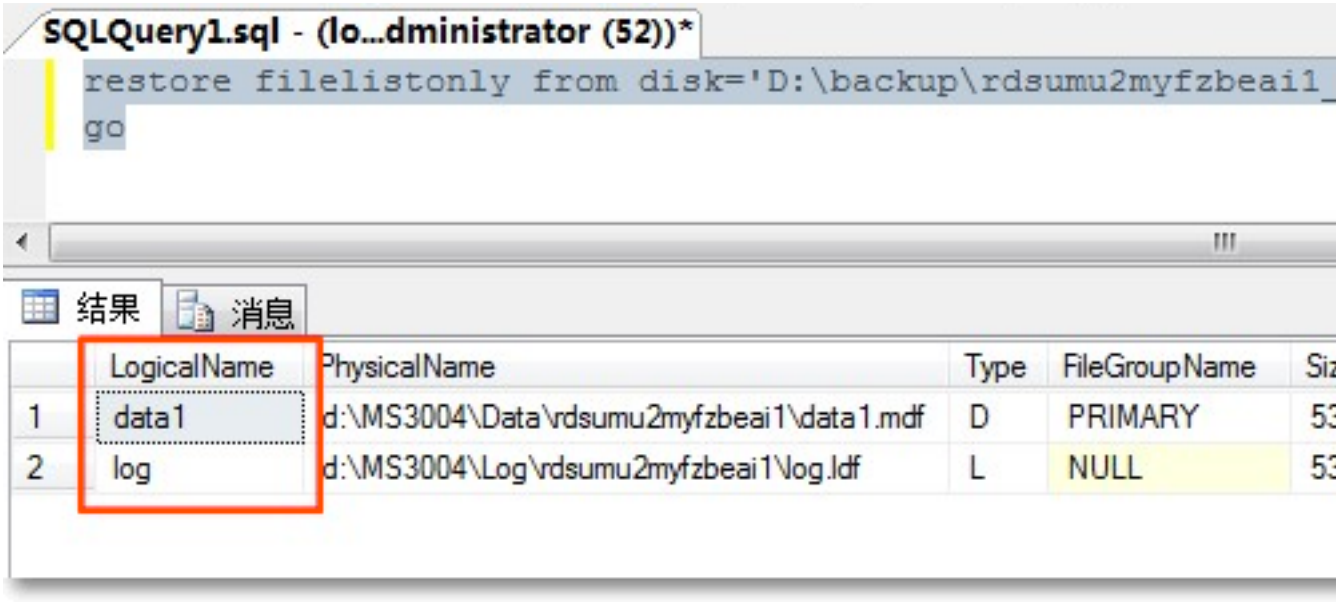

5. Load the full backup file.

```
restore database rdsumu2myfzbeai1 from disk=&apos;d:\backup\rdsumu2myfzbeai1_datafull_201402250050_2250050.bak&apos; with replace,norecovery,stats=10,
move &apos;data1&apos; to &apos;d:\database\rdsumu2myfzbeai1\data\data1.mdf&apos;,
move &apos;log&apos; to &apos;d:\database\rdsumu2myfzbeai1\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, the 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 rdsumu2myfzbeai1 from disk=&apos;D:\backup\
rdsumu2myfzbeai1_datadiff_201402260050_2260050.bak&apos; with
replace,norecovery,stats=10,
move &apos;data1&apos; to &apos;d:\database\rdsumu2myfzbeai1\data\
data1.mdf&apos;,
move &apos;log&apos; to &apos;d:\database\rdsumu2myfzbeai1\log\log.
ldf&apos;
go
```

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

7. Recover the database.

```
restore database rdsumu2myfzbeai1 with recovery
go
```

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

15.3.4 Migrate RDS for PostgreSQL data to the local PostgreSQL database

RDS for PostgreSQL supports the migration of cloud data to the local database by using logical backup files.

Procedure

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

```
pg_dump -U username -h hostname -p port databasename -f filename
```

Parameters description:

- username: Indicates the username used for database logon.

- hostname: Indicates the host name of the database.
- port: Indicates the database port number.
- databasename: 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 rds2z2tp80v3752wb455.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 desintationdb -p port -f dumpfilena
me.sql
```

Parameter description:

- username: Indicates the username used for database logon.
- hostname: Indicates the database address.
- port: Indicates the database port number.
- databasename: 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 the data import. They can be ignored, for example:

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

15.3.5 Migrate RDS for PPAS to local PPAS

ApsaraDB for PPAS supports the migration of cloud data to the local database by using logical backup files.

Procedure

1. Connect the PostgreSQL client to RDS.

2. Run the following command to back up the data.

```
pg_dump -U username -h hostname -p port databasename -f filename
```

Parameter descriptions:

- username: Indicates the username used for database logon.
- hostname: Indicates the host name of the database.
- port: Indicates the database port number.
- databasename: 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 rdsv07z563m7o25cj550public.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 desintationdb -p port -f dumpfilena
me.sql
```

Parameter descriptions:

- username: Indicates the username used for database logon.
- hostname: Indicates the database address.
- port: Indicates the database port number.
- databasename: 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 the data import. They can be ignored, for example:

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

15.4 Compress data with TokuDB for MySQL 5.6

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 lower. The TokuDB storage engine also 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 reading of data is 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. For MySQL 5.1 or 5.5, you have to upgrade it to MySQL 5.6 first.

2. Set the loose_tokudb_buffer_pool_ratio to indicate the proportion that TokuDB occupies in the shared cache of TokuDB and InnoDB.

```
select sum(data_length) into @all_size from information_schema.
tables where engine='innodb';
select sum(data_length) into @change_size from information_schema
.tables where engine='innodb' and concat(table_schema, '.',
table_name) in ('XX.XXXX', 'XX.XXXX', 'XX.XXXX');
select round(@change_size/@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.

15.5 Use psql to migrate PostgreSQL data

This example describes how to use the psql command 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 the backup file and then import it to the RDS through psql.

Prerequisites

You have set the whitelist, apply for an Internet address, and create 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 the data.

   ```
   pg_dump -U username -h hostname -p port databasename -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.
   - databasename: 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
   ```

Perform the migration

Note:

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

1. Log on to the ECS.
2. Run the following command through the PostgreSQL client to import the data into the RDS.

```
psql -U username -h hostname -d desintationdb -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
- databasename: 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 the data import. They can be ignored, for example:

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

15.6 Migrate SQL Server to cloud

15.6.1 Migrate data to ApsaraDB for RDS 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 the complete 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 the data migration to the cloud database.

Prerequisite

The migration target database is 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.

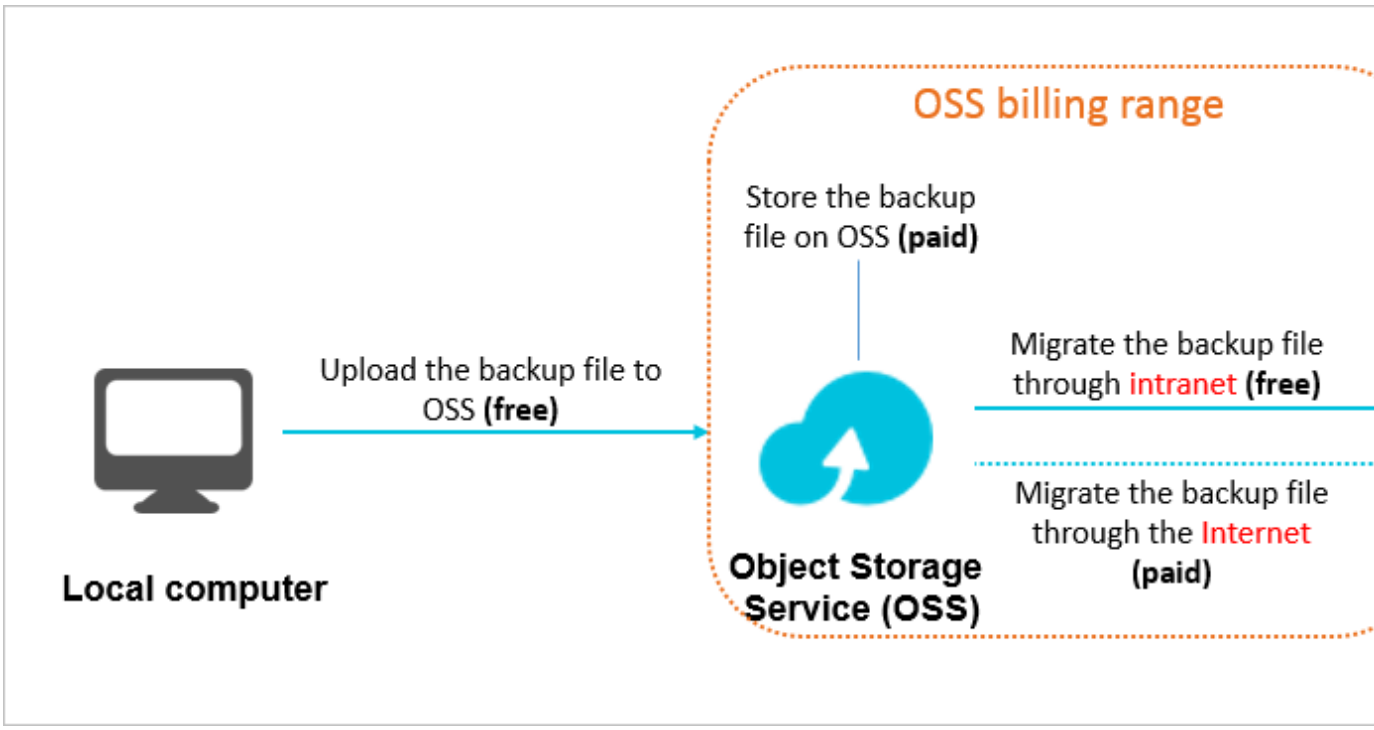

Figure description:

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

   c. Run the following command to check the Recover Mode of the local database.

   ```
   use master;
   go
   select name, case recovery_model
   when 1 then FULL
   when 2 then BULD_LOGGED
   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, perform Step d.
   - If the model value is FULL, perform Step e.

   d. Run the following command to set the Recover Mode of the source database to FULL.

   Note:
   Setting Recover Mode to FULL increases SQL Server logs. Therefore, make sure to leave 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 command 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
   COMPRESSION,INIT;
   go
   ```

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

   ```
   USE master
   GO
   RESTORE FILELISTONLY
     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. Back up the database again.

g. Run the following command 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 use a GUI, see *ossbrowser*.

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

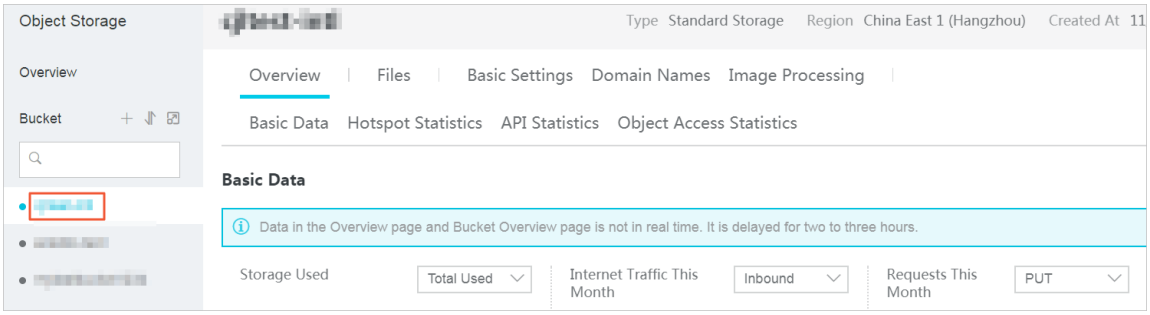

c. Select **Files**.

d. Click the name of the target backup file.

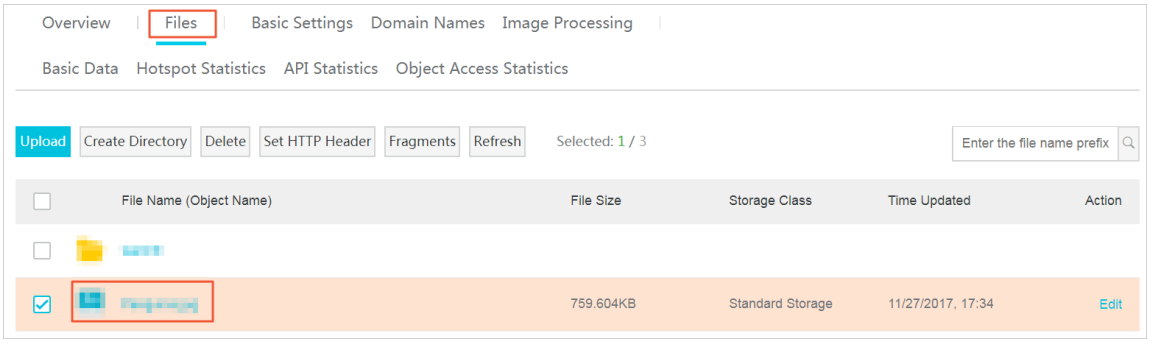

e. In the **Signature** field, change the validity period of the link. We recommend that you set the validity period to 28,800s, that is, eight hours.

> **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.

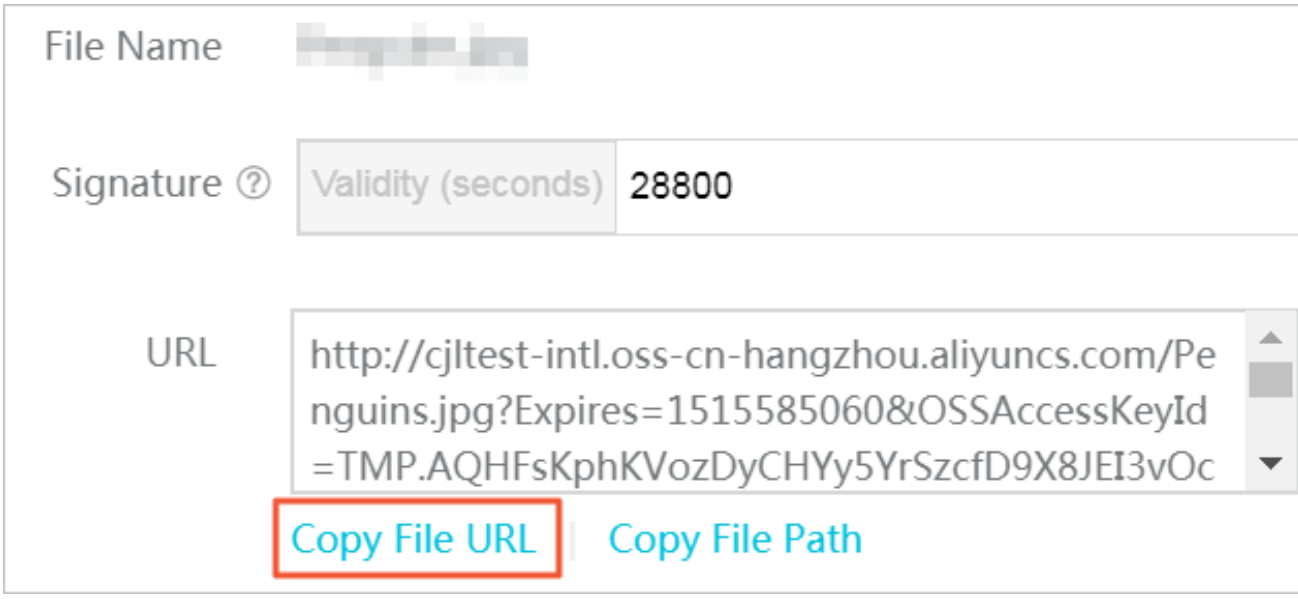

g. If you want to migrate data through the intranet, change the endpoint in the backup file URL to the intranet endpoint. The intranet endpoint varies depending on the network type and region. For more information, see *Access domain name and data center*.

For example, if the backup file URL is `http://rdstest-yanhua.oss-cn-shanghai.aliyuncs.com/testmigraterds_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, select **Databases** to go to the **Databases** page.

e. Find the target database and click **Migrate backup files from OSS** in the **Action** column.

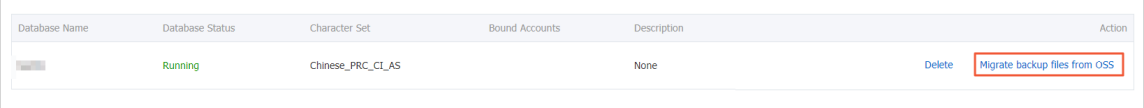

f. In the **Import Guide** dialog box, read the prompt and click **Next** to go to the **Upload the backup files** page.

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**.

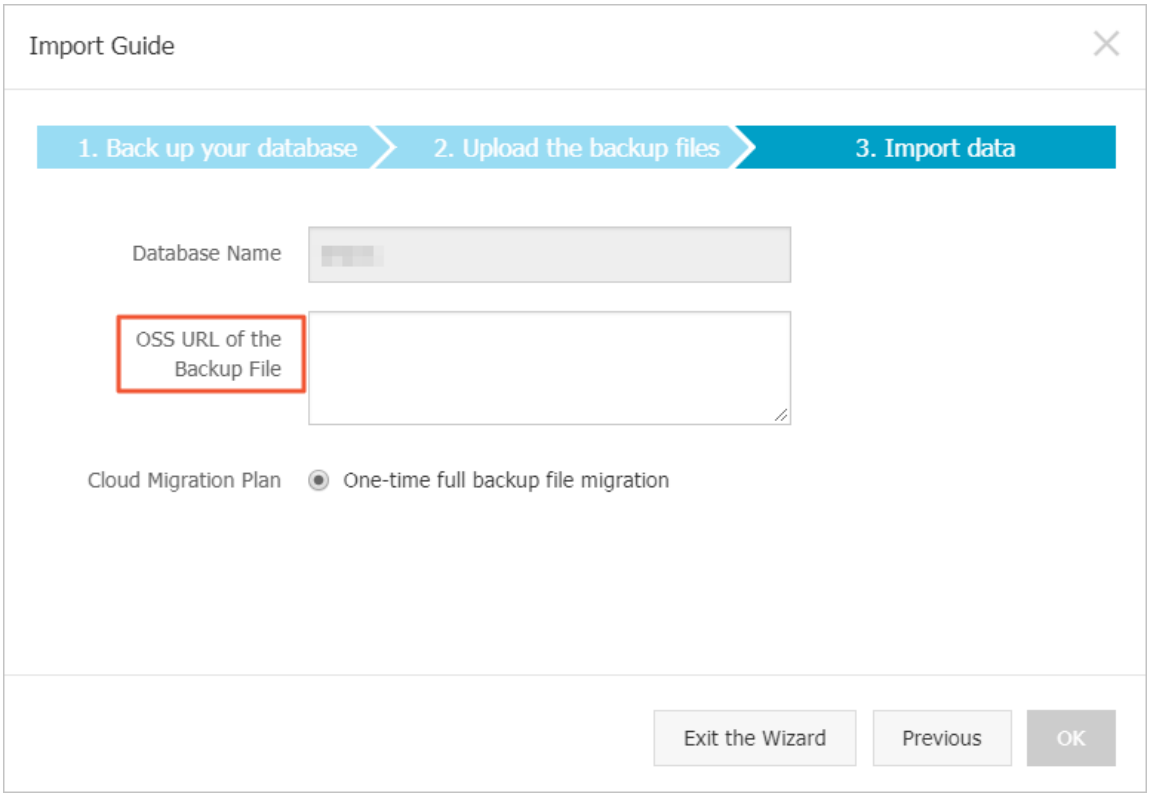

i. Click **OK**.

j. In the left-side navigation pane, select **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 the **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.

15.6.2 Migrate data to ApsaraDB for RDS SQL Server 2012/2016

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

Applicable versions

- Basic series (single-node): RDS SQL Server 2012 Web, Enterprise; RDS SQL Server 2016 Web, Enterprise
- High-availability series (dual-node): RDS SQL Server 2012 Standard, Enterprise; RDS SQL Server 2016 Standard, Enterprise

For instructions on how to migrate data to RDS SQL Server 2008 R2 Enterprise (high-availability series), see *#unique_150*.

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 the backup file is not generated using the script provided in this article, 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 full backup file name 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.

Backup file name

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

Delete backup file from OSS

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

Prerequisites

Instance capacity

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

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

You do not need to create a target database in advance. This is different from the requirement stated in *#unique_150*.

If a database with the same name already exists in the target instance, back up and delete the database before creating the migration task.

Create initial account on target instance

It is recommended that you create an initial account for the target instance on the console in advance. If the target instance does not have an initial account, the migration also succeeds but you cannot access the database unless you take measures by referring to Common Errors at the end of this article.

For information on how to create an initial account, see*########SQL Server 2012#* steps 1 to 7 in Create account and database for SQL Server 2012 and 2016.

Prepare 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 SQL Server instance so that the bucket can be selected in subsequent steps.)

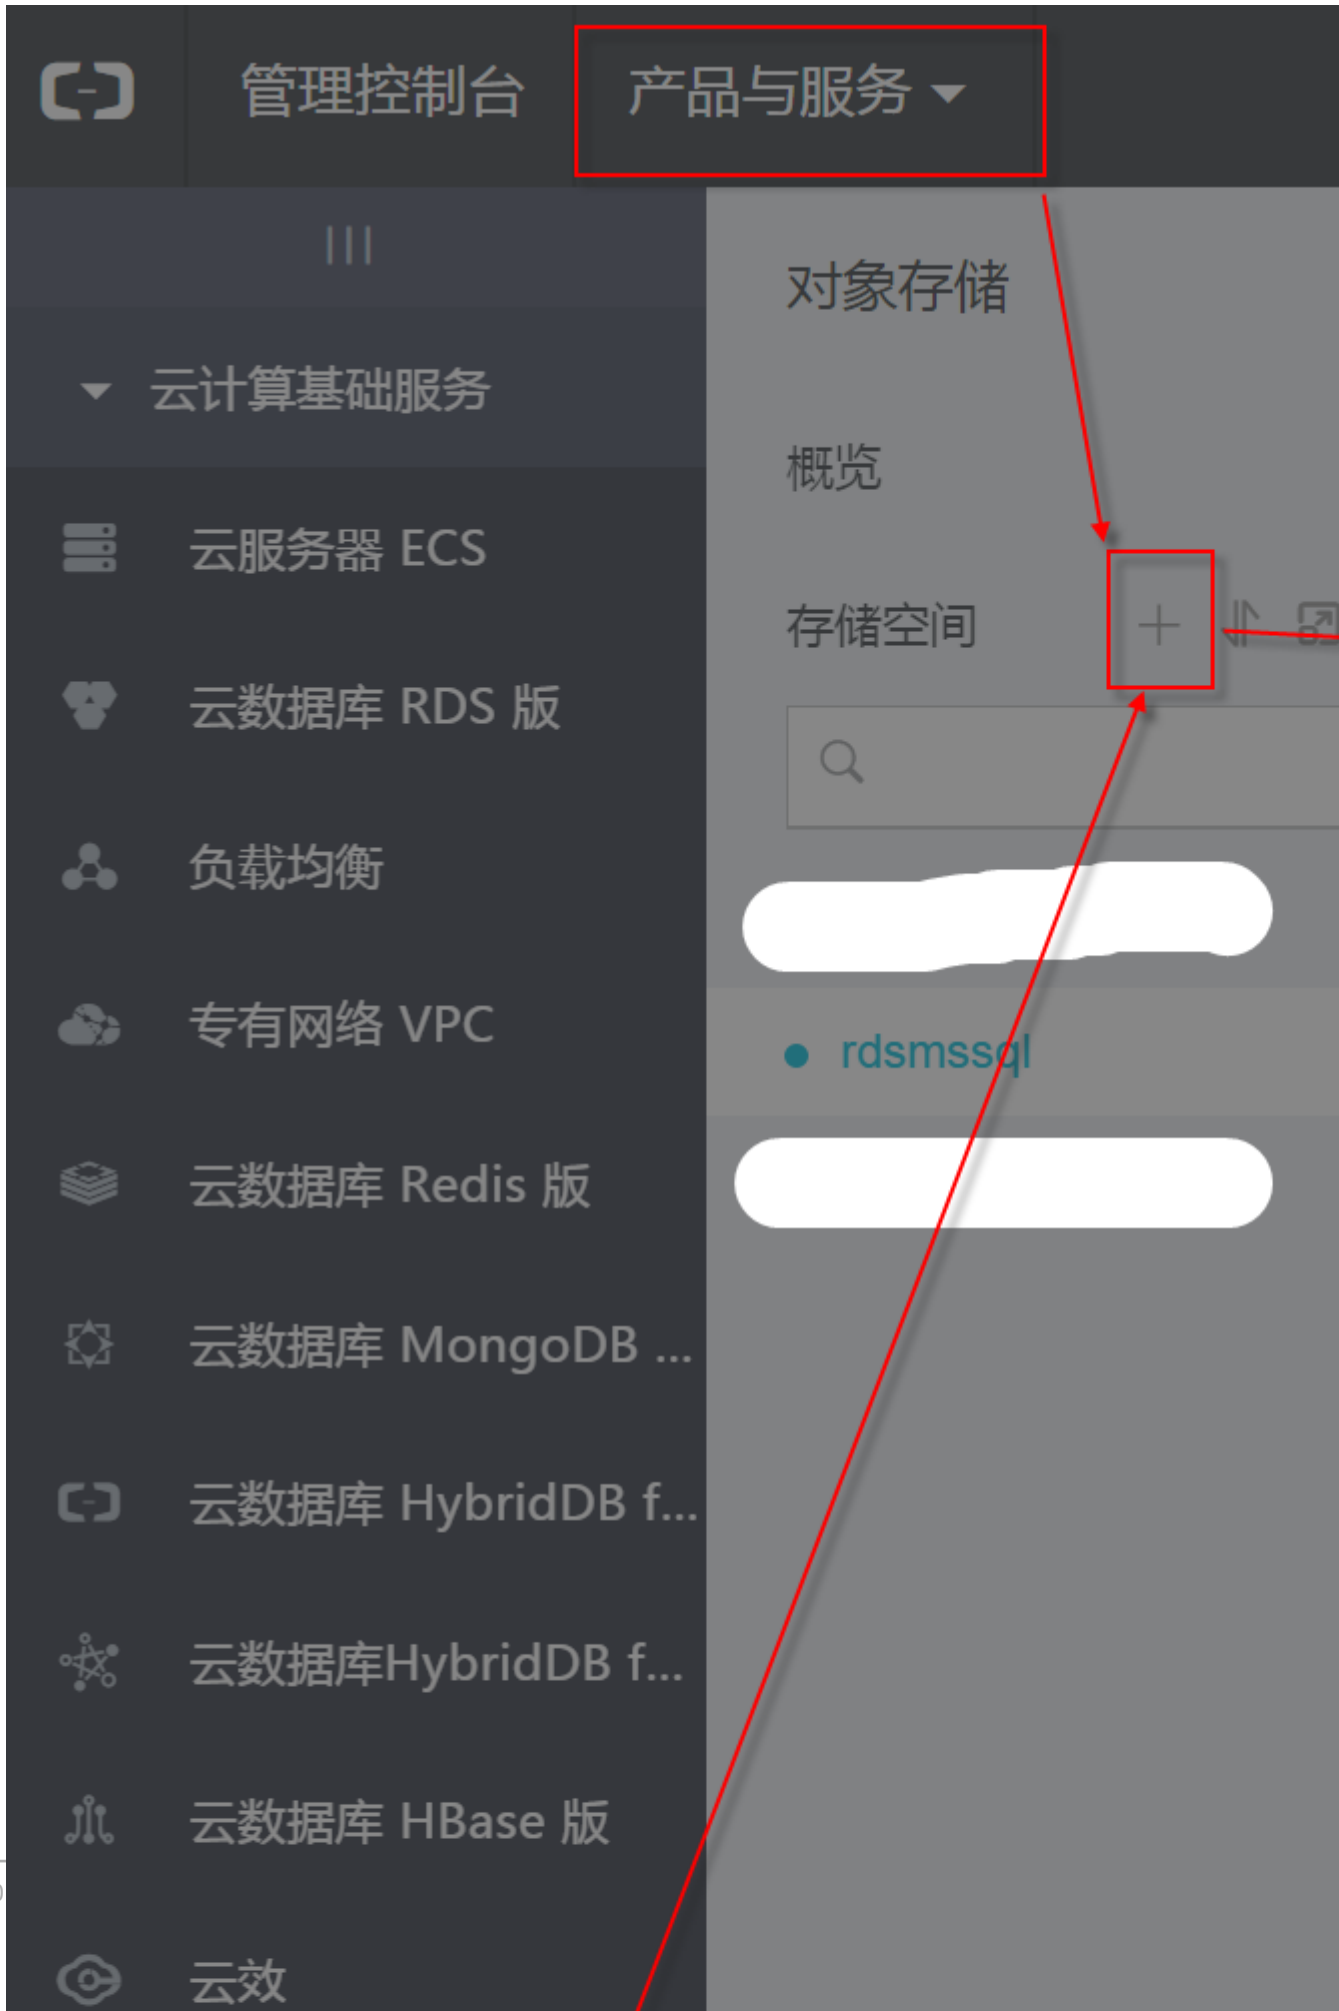
管理控制台
产品与服务
对象存储
概览
存储空间
rdsmssql
云计算基础服务
云服务器 ECS
云数据库 RDS 版
负载均衡
专有网络 VPC
云数据库 Redis 版
云数据库 MongoDB ...
云数据库 HybridDB f...
云数据库HybridDB f...
云数据库 HBase 版
云效

Run DBCC CHECKDB

Run DBCC CHECKDB('xxx') on the local database and ensure that the result has no allocation errors or consistency errors. The normal result is as follows:

```
CHECKDB found 0 allocation errors and 0 consistency errors in database
 'xxx'.
DBCC execution completed. If DBCC printed error messages, contact your
 system administrator.
```

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

Procedure

Only three steps are required to migrate a local database to RDS SQL Server 2012/2016 on the cloud:

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

Back up local databse

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 in a way you are used to or by following these steps:

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

Configuration item	Description
@backup_databases_list	databases to be backed up. Separate multiple databases with semicolon or comma.
@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 automatically created if it does not exist.
@is_run	whether to perform a backup. Values are as follows: The parameter values are as follows:

Configuration item	Description
	• 1: Perform a backup. • 0: Perform checking only.

3. Run the backup script.

Upload backup file to OSS

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

Create 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 access the **Basic Information** page.
4. On the left-side navigation pane, choose **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 offical service account to access OSS:

a. Click Authorize in the **data import** wizard:

b. Click **Confirm Authorization Policy**.

7. In **step 3** of the **data import wizard**, set the following parameters and click **OK** to generate the OSS backup file upload task.

Configuration item	Description
Database Name	target database name on the target instance
OSS Bucket	OSS bucket that stores the backup file
OSS Subfolder Name	Inlcude bucket subfolders if any. Separate subfolders of different levels with a slash (/), for example, Dir1/Dir2/Dir3. Skip this field if it is not involved.
OSS File	Click the magnifier icon on the right. You can perform a fuzzy search of the backup file prefix. The file names, sizes, and update time

Configuration item	Description
	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 differenti al or log backup file, select this option.
Consistency Check Mode	• Synchronous DBCC: This option performs DBCC Check only after the database is opened. This option 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 recify faults by referring to Common errors at the end of this article.

View migration records

View migration records as follows:

On the ackup 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 failuer 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: A existing database with the same name is not allowed on the target instance. This prevents you from mistakenly overwriting a database.
- Solution: If a database with the same name already exists on the target instance, perform a full backup of the database on the console and delete the database before the migration.

Differential backup file

- 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 file

- Error message: Backup set (xxx.trn) is a Transaction Log backup, we only accept a FULL Backup.
- Error cause: The 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 if 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 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 error

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

  1. Use the following command to fix the local database (Note: This may cause data loss):

     ```
     DBCC CHECKDB (DBName, REPAIR_ALLOW_DATA_LOSS) WITH NO_INFOMSGS,
     ALL_ERRORMSGS
     ```

  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 SQL 2008 R2 high-availability edition.

- 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). See the following figure.


对象存储
概览
存储空间
rdsmssql
概览
文件
上传文件
新建目
文件名
testmig

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

预览

文件名 testmigraterds_20170906143807_

签名 ⓘ 链接有效时间（秒） 64800

URL http://rdsmssql.oss-cn-hangzhou
aterds_20170906143807_FULL.ba
7&OSSAccessKeyId=TMP.AQFn7

复制文件 URL | 复制文件路径

类型 application/octet-stream

> **Note:**
> Note: The backup file with the Public Read ACL permission is always downloadable without an expiration date of the download link. To prevent security risks, set the ACL permission to **Private** after migrating the file.

Insufficent space 1

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

Insufficent space 2

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

No initial 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 initial account, the migration still succeeds, but the migration task does not know which user to authorize.
- Solution:

  1. Create an initial account. For details, see steps 1 to 7 in *########SQL Server 2012#*.
  2. Reset the password of the initial account. For more information, see *Reset instance password*.
  3. Use the initial account to access the database on the cloud.

End

序号	数据库名	任务开始时间	任务结束时间	任务状态	任务类型	任务描述
100674	testdb	2018-02-09 21:53:06	2018-02-09 21:59:55	成功	全量备份文件一次性迁入	success
100673	testdb3	2018-02-09 21:31:07	2018-02-09 21:46:45	失败	全量备份文件一次性迁入	Failed to download backup file
100672	adventureworks2008r2	2018-02-09 21:30:02	2018-02-09 21:45:40	失败	全量备份文件一次性迁入	Your backup is corrupted or n
100671	testdb2	2018-02-09 21:29:16	2018-02-09 21:43:00	失败	全量备份文件一次性迁入	DBCC checkdb failed
100670	testmigrate	2018-02-09 21:28:58	2018-02-09 21:37:35	成功	全量备份文件一次性迁入	success
100669	testdb	2018-02-09 21:28:39	2018-02-09 21:32:05	失败	全量备份文件一次性迁入	Your backup is corrupted or n
100668	testdb3	2018-02-09 21:28:20	2018-02-09 21:30:55	失败	全量备份文件一次性迁入	autotest_2008r2_std_testmigr only accept a FULL Backup.
100667	adventureworks2008r2	2018-02-09 21:28:02	2018-02-09 21:29:30	失败	全量备份文件一次性迁入	autotest_2008r2_std_testmigr backup, we only accept a FUL

OSS下载URL有效期过期，上云失败

DBCC Checkdb失败，导致上云失败

Simple read_only数据库备份文件，上云成功

日志备份，上云失败

差异备份，上云失败

16 Typical applications

16.1 Cached data persistence

RDS can be used together with ApsaraDB for Memcache and ApsaraDB for Redis to form a storage solution with high throughput and low delay. The following section describes the cached data persistence solution based on the combined use of RDS and ApsaraDB for Memcache.

Background information

Compared with the RDS, the RDS cache product has the following two features:

- High response speed: The request delay of the RDS for Memcache and the RDS for Redis is usually within several milliseconds.
- The cache area can support a higher QPS (Requests Per Second) than the RDS.

System requirements

- Bmemcached (with support of SASL extension) has been installed in the local environment or ECS.

  Bmemcached download address: Click 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 RDS and ApsaraDB for Memcache:

```
/usr/bin/env python
import bmemcached
Memcache_client = bmemcached.Client((‘ip:port’), ‘user’, ‘passwd’)
#Search for a value in ApsaraDB for Memcache
res = os.client.get(‘test’)
if res is not None:
    return res #Return the searched value
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

The OSS is a cloud storage service provided by Alibaba Cloud, featuring massive capacity, security, low cost, and high reliability. The RDS can work with the OSS to form multiple types of data storage solutions.

For example, when the business application is a forum and the RDS works with the OSS, resources such as registered users' images and post content images can be stored in the OSS to reduce the storage pressure of the RDS.

Sample code

The OSS works with the RDS example.

1. Initialize OssAPI.

```
from oss.oss_api import *
endpoint=”oss-cn-hangzhou.aliyuncs.com”
accessKeyId, accessKeySecret=”your id”,”your secret”
oss = OssAPI(endpoint, accessKeyId, accessKeySecret)
```

2. Create a bucket.

```
#Set the bucket to private-read-write
res = oss.create_bucket(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”
accessKeyId, accessKeySecret=”your id”,”your secret”
oss = OssAPI(endpoint, accessKeyId, accessKeySecret)
User_id = mysql_client.fetch_one (SQL) # Search for user_id in RDS
#Obtain and download the user avatar to the corresponding 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’)
```

17 Appendix

17.1 Commonly used SQL commands for MySQL

This document lists some of the commonly used SQL commands. Only the sytaxes are explained. For the detailed information on SQL commands, including command parameters and restrictions, see *MySQL 5.7 Reference Manual*.

Database-related commands

Operation	Command
Create a database and designate a character set	`create database db01 DEFAULT CHARACTER SET gbk COLLATE gbk_chinese_ci;`
Delete a database	`drop database db01;`

Account-related commands

Note:

If an instance has the high-privilege account, the passwords of other accounts/users under this instance cannot be changed through the high-privilege account. If the password needs to be changed, you must delete this account and create a new one.

Operation	Command
Create an account	`CREATE USER 'username'@'host' IDENTIFIED BY 'password';`
Delete an account	`DROP USER 'username'@'host';`
Authorization	`GRANT SELECT ON db01. * TO 'username'@'host';`

Query the created accounts in the database	`SELECT user,host,password FROM mysql.user_view;`
	or
	`show grants for xxx`

Reclaim permissions	• Reclaim all permissions
	`REVOKE ALL PRIVILEGES,GRANT OPTION FROM 'username'@'host';`
	• Reclaim specific permissions
	`REVOKE UPDATE ON *. * FROM 'username'@'host';`

17.2 View instance intranet/Internet address and port number

When connecting to an RDS instance, you must enter the intrane/Internet address and port number of the target RDS instance. This document introduces where to view these information on 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 target instance to enter the **Basic Information** page.
4. In the **Basic Information** area, you can find the Internet/intranet address and Internet/intranet port number of the RDS instance, as shown in the following figure.


Basic Information
Set White List
Instance ID:
Name:
Instance Region and Zone: China East 1 (Hangzhou)ZoneB
Instance Type: Standard (rds.status.category.Basic)
Intranet Address
Intranet Port: 3306
Internet Address
Outer Port: 3306