

Alibaba Cloud ApsaraDB for MySQL

User Guide

Issue: 20181129

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 <code>cd /d C:/windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is a optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	It indicates that it is a required value, and only one item can be selected.	<code>swich {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Preface.....	1
2 Quick start.....	3
3 Billing management.....	5
3.1 Change the billing method.....	5
3.2 Enable auto-renewal for a Subscription instance.....	6
3.3 Manually renew a Subscription instance.....	7
4 Instance management.....	9
4.1 Restart an instance.....	9
4.2 Configure the maintenance period.....	9
4.3 Migrate instance across zones.....	10
4.4 Switch between master and slave instances.....	15
4.5 Modify the data replication mode.....	17
4.6 Create a read-only instance.....	19
4.7 Release an instance.....	21
4.8 Upgrade the database version.....	22
4.9 RDS for MySQL release notes.....	23
4.10 Change configurations.....	24
4.11 SQL Server DBCC function.....	27
4.12 End connections for SQL Server instances.....	28
5 Connection management.....	29
5.1 Set the access mode.....	29
5.2 Set network type.....	29
5.3 Hybrid access solution for smooth migration from classic networks to VPCs.....	31
5.4 Set intranet and Internet IP addresses.....	36
6 Monitoring and Alarming.....	41
6.1 Set the monitoring frequency.....	41
6.2 Set monitoring rules.....	46
7 Security.....	48
7.1 Switch the IP whitelist to enhanced security mode.....	48
7.2 Set SSL encryption.....	49
7.3 Set TDE.....	52
8 Log management.....	55
9 Backup.....	57
9.1 Back up RDS data.....	57
9.2 View the free quota of the backup space.....	61
9.3 Download data and log backup files.....	61

9.4 Logical backup and recovery for PPAS.....	63
10 Recovery.....	66
10.1 Create a clone instance.....	66
10.2 Recover data directly to an instance.....	69
11 Create a linked server for SQL Server instances.....	71
12 Data migration.....	73
12.1 Migrate from other cloud databases to ApsaraDB RDS.....	73
12.2 Use mysqldump to migrate MySQL data.....	73
12.3 Migrate RDS data to the local database.....	76
12.3.1 Migrate RDS for MySQL data to a local MySQL database.....	76
12.3.2 Migrate RDS for SQL Server data to a local SQL Server database.....	79
12.3.3 Migrate RDS for PostgreSQL data to a local PostgreSQL database.....	81
12.3.4 Migrate RDS for PPAS data to a local Oracle database.....	82
12.3.5 Migrate RDS for PPAS data to a local PPAS database.....	85
12.4 Compress data.....	86
12.5 Use psql to migrate PostgreSQL data.....	87
12.6 Migrate SQL Server to cloud.....	89
12.6.1 Migrate data to RDS for SQL Server 2008 R2.....	89
13 Typical applications.....	96
13.1 Cached data persistence.....	96
13.2 Multi-structure data storage.....	97

1 Preface

Overview

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

This document describes RDS features and functions and further explains the procedure to configure RDS through the [RDS console](#). You can also manage RDS through APIs and SDKs.

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

For more information about functions and pricing of RDS, log on to [official website of ApsaraDB for RDS](#).

Declaration

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

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

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

Consideration

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

General terms

- **Instance:** A database service process that takes up physical memory independently. You can set different memory size, disk space, and database type, among which the memory specification determines the performance of the instance. After the instance is created, you can change the configuration and delete the instance at any time.

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

Common conventions

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

2 Quick start

If you use RDS for the first time, see the following *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 *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 and a combination of open source software (Linux + Apache + MySQL + Perl/PHP/Python), MySQL is widely used in a variety of applications.

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

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

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

ApsaraDB for SQL Server

SQL Server is one of the first commercial databases and is an important part of the Windows platform (IIS + .NET + SQL Server), with support for a wide range of enterprise applications. The SQL Server Management Studio software comes with a rich set of built-in graphical tools and script editors. You can quickly get started with a variety of database operations through visual interfaces.

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

RDS for SQL Server currently supports the following versions:

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

ApsaraDB for PostgreSQL

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

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

RDS for PostgreSQL currently supports version 9.4.

ApsaraDB for PPAS

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

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

RDS for PPAS currently supports version 9.3.

3 Billing management

3.1 Change the billing method

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

Attention

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



Note:

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

Prerequisites

- You are the owner of the instance.
- The instance type is not a history instance type. For more information, see [Instance type overview](#).



Note:

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

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


**Note:**

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

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

Procedure

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



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

3.2 Enable auto-renewal for a Subscription instance

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

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

Prerequisite

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

Attentions

- The renewal cycle cannot be changed while enabling the auto-renewal function. For variable renewal cycles, renew the instance manually. For more information about how to handle manual renewal, see [Manually renew a Subscription instance](#).

- If you select auto-renewal, you are charged three days before the instance expires. Credit cards and coupons are supported for each renewal payment.
- If you manually renew your instance before the charging date, auto-renewal takes place based on the new expiration date.
- The auto-renewal function takes effect the next day after it is enabled. If your instance expires on the next day, manually renew it to prevent service interruptions.

Procedure

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




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

3.3 Manually renew a Subscription instance

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

Procedure

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the ID of the target instance to go to the **Basic Information** page.
4. Click **Renew** in **Status** area, as shown in the following figure.



5. Select the renewal period on the **Renew** page.

**Note:**

You can change the configuration if needed.

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

Related topic

[*Enable auto-renewal of the subscription instance*](#)

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](#).
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 Configure the maintenance period

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

Background information

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


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

To guarantee stability during the maintenance process, instances enter the **Instance being maintained** state before the preset maintenance period on the day of maintenance. When an instance is in this state, normal data access to databases is not affected. However, apart from account management, database management, and IP address addition to the whitelist, other services involving changes (such as common operations including upgrade, degrade, and restart of the instance) are unavailable on the console. Query services such as performance monitoring are available.

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

Procedure

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




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



Note:

Note: The time segment is in Beijing Time.



4.3 Migrate instance across zones

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



Note:

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

Background information

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

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

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

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

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

**Note:**

Because certain network delay exists between multi-zones, the response time of a multi-zone instance to a single update may be longer than that of a single-zone instance when a multi-zone

instance adopts the semi-synchronous data replication mode. In this case, increase the overall throughput by enhancing the concurrency.

Attentions


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

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


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

Procedure

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




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



Parameter description:


- **Migrate to:** Select the region to which you want to migrate the instance.

- **Switching Time:** Choose when to perform the migration. During the migration, many operations cannot be performed. You can choose to switch immediately or at a later time.
6. To modify the maintenance time, perform the following steps. Alternatively, you can also leave the maintenance time unchanged.
- a. Click **Modify**, as shown in the following figure. The **Basic Information** page is displayed.



Switching Time: ☒ **Switch now**
☐ **Switch within maintenance period**(current setting : 02:00-06:00 [Modify](#))

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



Configuration Information

Class Family: General

Database Memory: 1024MB

Maintenance Period:

☐ 06:00-07:00	☐ 07:00-08:00	☐ 08:00-09:00
☐ 09:00-10:00	☐ 10:00-11:00	☐ 11:00-12:00
☐ 12:00-13:00	☐ 13:00-14:00	☐ 14:00-15:00
☐ 15:00-16:00	☐ 16:00-17:00	☐ 17:00-18:00
☐ 18:00-19:00	☐ 19:00-20:00	☐ 20:00-21:00
☐ 21:00-22:00	☐ 22:00-23:00	☐ 23:00-00:00
☐ 00:00-01:00	☐ 01:00-02:00	☐ 02:00-03:00
☐ 03:00-04:00	☐ 04:00-05:00	☐ 05:00-06:00

[Save](#) [Cancel](#)

- c. Go back to the page for migrating the instance to another zone.
7. In the **Migrate Instance to Other Availability Zones** dialog box, click **OK**.

Migrate Instance to Other Availability Zones

Node 's Instance:

Current Availability Zone: ZoneB

Migrate to:

China East 1 (Hangzhou)ZoneD

China East 1 (Hangzhou)ZoneG

Current VPC: vpc-

No virtual switch exists in the VPC of current zone. [please create a new switch first on the VPC console.](#)

Switching Time:

☐ Switch now
☒ Switch within maintenance period(current setting : 02:00-06:00 [Modify](#))

You will be disconnected for 30s during the availability zone migration. To proceed with migration, the database must have a reconnection mechanism.

OK

Cancel

4.4 Switch between master and slave instances

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

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

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

Attentions

- Currently this operation is not applicable to the Basic Edition of MySQL 5.7 and SQL Server 2012/2016 instances. This is because Basic Edition instances do not have slave nodes.
- Switching between master and slave instances may result in transient 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 switch, many operations cannot be performed. Therefore, we recommend that you choose to switch within the maintenance period.

Master/Slave Node Switchover

Are you sure you want to proceed with master/slave node switch? You may experience 1 or 2 disconnections.

Switching Time:

☒ **Switch now**

☐ **Switch within maintenance period** (current setting : 02:00-06:00 [Modify](#))

6. Do as follows to change the maintenance period is necessary:
 - a. Click **Modify** to open the **Basic Information** page.

Switching Time:

☒ **Switch now**

☐ **Switch within maintenance period** (current setting : 02:00-06:00 [Modify](#))

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

Configuration Information

Class Family: General

Database Memory: 4096MB

Maintenance Period:

☐ 06:00-07:00	☐ 07:00-08:00	☐ 08:00-09:00
☐ 09:00-10:00	☐ 10:00-11:00	☐ 11:00-12:00
☐ 12:00-13:00	☐ 13:00-14:00	☐ 14:00-15:00
☐ 15:00-16:00	☐ 16:00-17:00	☐ 17:00-18:00
☐ 18:00-19:00	☐ 19:00-20:00	☐ 20:00-21:00
☐ 21:00-22:00	☐ 22:00-23:00	☐ 23:00-00:00
☐ 00:00-01:00	☐ 01:00-02:00	☐ 02:00-03:00
☐ 03:00-04:00	☐ 04:00-05:00	☐ 05:00-06:00

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

7. Click **OK**.

4.5 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:

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

Background information


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

- Sync mode:


- When the updates initiated by applications are all executed at the master node, logs are synchronously transferred to all the slave nodes. The transaction completes the commit only when most nodes (including the master node) in the cluster have received and stored the logs.
- Only instances that have three or more nodes support strong synchronous replication. No matter what happens, the strong synchronous replication mode cannot be degraded to 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 applications until the replication mode times out and degrades to the async mode.
 - When data replication between the two nodes resumes normally (the slave node or network connection is recovered), async mode is changed to sync mode. The time period required for restoration to the sync mode depends on the implementation mode of the semi-sync mode. ApsaraDB for MySQL 5.5 differs from ApsaraDB for MySQL 5.6 in this regard.
- Async mode: An application initiates an update (including addition, deletion, and modification operations) request. After completing the corresponding operation, the master node immediately responds to the application and then replicates data to the slave node asynchronously. Therefore, in the async mode, unavailability of the slave node does not affect the operation on the slave database, and unavailability of the master node has a low probability to cause data inconsistency between the two nodes.

Procedure

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the ID of the target instance to visit the **Basic Information** page.
4. In the left-side navigation pane, select **Instance Availability**.
5. Click **Modify Data Replication Mode**, as shown in the following figure.



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



7. Click **OK**.

4.6 Create a read-only instance

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

Attention

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

— SQL Server 2017


- Quantity of read-only instances

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

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

Procedure

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



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



Note:

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

6. Select **Product Terms of Service and Service Level Notice and Terms of Use**, and then click **Pay Now**.

7. After creating the read-only instance, you can view it on the **Instances** page, as shown in the following figure.

☐ Instance Name	Status(All) ▼	Creation Time	Instance Type(All) ▼
☐  [blurred instance name]	Running	2017-07-26 16:24	Read-only
☐  [blurred instance name]	Running	2017-07-18 15:03	Regular

4.7 Release an instance

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

Attentions

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

Procedure

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

Operating Status			Release Instance	^
Operating Status: Running	Billing Method: Pay-As-You-Go	Creation Time: 2017-07-18 15:03:03		

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

4.8 Upgrade the database version

Background information

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

Attentions

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

Procedure

- Log on to the [RDS console](#).
- Select the region where the target instance is located.
- Click the ID of the target instance to enter the **Basic Information** page.
- In the **Configuration Information** area, click **Upgrade Database**, as shown in the following figure.

Configuration Information

Class Family: General

Database Engine: MySQL 5.5 Upgrade Database

CPU: 1Core

Database Memory: 1024MB

Maximum IOPS: 600

Maximum Number of Connections: 300

Time Segment: 02:00-06:00 Settings

Instance Class: rds.mysql.t1.small

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

4.9 RDS for MySQL release notes

MySQL 5.7

mysql57_20180431:

- New features:
 - Supports the High-availability Edition.
 - Supports the [database proxy](#) function.
 - Supports [SQL audit](#).
 - Enhanced protection for instances that are generating snapshots.

MySQL 5.6

- **mysql_201806** (5.6.16) (coming soon):**
 - New feature: Increases the slow log precision to microsecond.
- **mysql_20180426 (5.6.16)**
 - New feature: Supports hidden indexes so that you can set invisible indexes. For more information, see [Reference](#).
 - Bugs fixed:
 - Fixed bugs that occur when backup instances are applying threads.
 - Resolved the performance deterioration that occurs when backup instances are applying partition updates.
 - Resolved the problem that an entire TokuDB table is rebuilt by the ALTER TABLE COMMENT command. For more information, see [Reference](#).
 - Resolved possible deadlocks triggered by the SHOW SLAVE STATUS or SHOW STATUS command.
- **mysql_20171205 (5.6.16):**
 - Resolved the problem that concurrent execution of OPTIMIZE TABLE and ONLINE ALTER TABLE causes deadlocks.
 - Resolved conflicts between SEQUENCE and implicit primary keys.
 - Resolved problems related to SHOW CREATE SEQUENCE.
 - Resolved the problem that TokuDB table statistics are incorrect.
 - Resolved the problem that parallel OPTIMIZE table commands cause deadlocks.
 - Resolved the character set problems recorded in QUERY_LOG_EVENT.

- Resolved the problem that databases cannot be stopped due to signal processing. For more information, see [Reference](#).
- Resolved problems caused by RESET MASTER.
- Resolved the problem that backup databases are stuck in the waiting state.
- Resolved the status maintenance problem caused by master node failovers of Finance Edition instances.
- Resolved the possible process termination caused by SHOW CREATE TABLE.
- **mysql_20170927 (5.6.16):**
 - Resolved the problem that TokuDB table queries use incorrect indexes.
- **mysql_20170901 (5.6.16):**
 - New features:
 - The SSL encryption version is upgraded to TLS1.2. For more information, see [Reference](#).
 - SEQUENCE is supported.
 - Resolved the problem that NOT IN queries return incorrect results in certain scenarios.
- **mysql_20170530 (5.6.16):**
 - New feature: A master account can kill connections of common accounts.
- **mysql_20170221 (5.6.16):**
 - New feature: [Read/write splitting](#) is supported

4.10 Change configurations

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

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

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

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

Attention

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

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

Function	Supported or not
Change Backup Mode	No
Restore Data	No
Modify Parameters	No

Procedure

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the **Basic Information** page.
4. In the configuration information bar, click **Change configuration** to go to the change instance page.
5. In the change configuration bar, select a new configuration.

Parameter description:

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




Note:

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


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

6. Do as follows if you want to modify the maintenance period. Otherwise, skip the steps.

- a. Click **Modify**, as shown in the following figure. The system opens a new page and turns to the **Basic Information** page of the instance.



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



- c. Returns to the page for changing instance configurations.

7. On the instance configuration change page, click **Confirm**. For Subscription instances, complete the payment process according to subsequent prompts.

4.11 SQL Server DBCC function

RDS for SQL Server 2012 and later versions supports some features related to Database Console Commands (DBCC). You only need to use the stored procedure `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 commands:

```
USE master
GO
--database engine edition
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
```

4.12 End connections for SQL Server instances



Note:

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

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

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

5 Connection management

5.1 Set the access mode

This function has been replaced by the database proxy function. For more information, see [Database proxy](#).

5.2 Set network type

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

**Note:**

To migrate an instance from a classic network to a VPC without service interruptions, see [Hybrid access solution for smooth migration from classic networks to VPCs](#).

Background information

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

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

Precautions

- After switching the network type, the original intranet IP address is changed and the Internet IP address remains unchanged. Update the connection address on your applications if necessary. For example, after an RDS instance is switched from a classic network to a VPC, the intranet IP address of the classic network is released and a VPC IP address is generated. Therefore, ECS instances in classic networks cannot access the RDS instance through the intranet any more.
- To switch MySQL 5.5, MySQL 5.6, or SQL Server 2008 R2 instances from a classic network to a VPC, the access mode must be set to safe connection mode. To switch the access mode, see [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 network type switching, RDS services may be interrupted for about 30 seconds. Therefore, switch the network type during off-peak hours or make sure that your applications have the automatic reconnection mechanism.

Procedure

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

**Note:**

- If the drop-down lists do not display VPCs or virtual switches or if the VPCs and virtual switches are not what you need, create a VPC and virtual switch that are in the same region as the RDS instance. To create a VPC, see [Create a VPC](#). To create a virtual switch, see [Create a switch](#).
- For MySQL 5.5, MySQL 5.6, and SQL Server 2008 instances, their access mode must be safe connection mode if you want to switch from a classic network to a VPC. To switch the access mode, see [Set access mode](#).



3. Click **OK**.

- Switch from a VPC to a classic network

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

2. Click **OK**.

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

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

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

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


Background information

With a traditional solution, migrating an RDS instance from a classic network to a VPC causes immediate release of classic network IP address. As a result, an intermittent interruption for up to 30 seconds may be caused, and ECS on the classic network can no longer access the RDS

instance using the intranet IP address, which may have negative impact on your services. In many large companies, a database is usually designed for access by more than one application system. When they decide to migrate the database from a classic network to a VPC, it would be quite difficult to migrate the network of all the applications simultaneously, which may result in bigger impact on their services. Therefore, a transitional period is required. To accommodate the need for smooth migration, RDS has added the hybrid access feature, making it possible to have such a transitional period.

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

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



Functional Limits

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

- Switch to classic networks is not supported.
- Zone migration is not supported.
- 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 [Database proxy overview](#). 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, create them by referring to [Create VPC](#) and [Create VSwitch](#).

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



Note:

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

Switch to VPC

!

Switch to classic, include endpoint(s):
Intranet Port: [Set Whitelist](#) and then address will be shown.

Switch to:
VPC: vpc- Virtual Switch:vsw-

If the switch you need is not in the list, [please create a new switch first on the VPC console.](#)

Note: Switching to Virtual Private Cloud (VPC) will cause an intermittent interruption, and the ECS in the classic network will not be able to access the database. If you need to reserve the Intranet address of the classic network, check the following option.

☒ **Reserve original classic endpoint**

The hybrid access solution reserves the Intranet address of the original classic network and adds an Intranet address under VPC, which prevents any intermittent interruption during migration and has no impacts on your service. You are advised to use VPC only for the sake of security and performance. For this reason, the reserved Intranet address of the classic network is available for a limited period of time and will be released once the reserved period expires. In that case, your applications will not be able to access the database using the Intranet address of the classic network.

Expiration time

14 day(s) 30 day(s) 60 day(s) 120 day(s)

OK

Cancel

8. Click **OK**.

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

Issue: 20181129

35

Change the expiration time of the original classic network

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

Follow these steps to change the expiration time:

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

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

5.4 Set intranet and Internet IP addresses

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

Background information

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

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

The applicable scenarios of the connection addresses are as follows:

- Use the intranet IP address only:
 - The system provides an intranet IP address by default and you can directly modify the connection address.
 - This scenario is applicable when your application is deployed on the ECS instance that is located in the same region and has the same network type as your RDS instance.
- Use the Internet IP address only:
 - This scenario is applicable when your application is deployed on the ECS instance that is located in the different region from that of your RDS instance.
 - This scenario is applicable when your application is deployed on a platform other than Alibaba Cloud.
- Use both of the intranet and Internet IP addresses:


- This scenario is applicable when your application is deployed on the ECS instance that is located in the same region and has the same [network type](#) as your RDS instance, and application modules are deployed in an ECS where your RDS instance is not located.
- This scenario is applicable when your application is deployed on the ECS instance that is located in the same region and has the same [network type](#) as your RDS instance, and on a platform other than Alibaba Cloud.

Attentions

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

Apply for an Internet IP address

1. Log on to the [RDS console](#) .
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the **Basic Information** page.
4. Click **Connection options** in the left-side navigation pane.
5. Click **Apply for Internet Address**, as shown in the following picture.




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

Modify the connection address

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

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

4. Click **Connection options** in the left-side navigation pane.
5. Click the **Instance Connection** tab.
6. In the **Connection 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 contains 8 to 64 characters including letters and digits. It must begin with a lower-case letter.
- **Port:** indicates the number of the port through which RDS provides external services, which can be an integer within the range [3200, 3999].

Release an Internet IP address


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



Note:

The operation can be performed only in **safe connection mode**. For more information about the safe connection mode, see [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. Click **Connection options** in the left-side navigation pane.
5. Click the **Instance Connection** tab.
6. In the **Connection Information** area, click **Release Internet Address**.



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

6 Monitoring and Alarming

6.1 Set the monitoring frequency

Background information

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

Two monitoring frequencies provided before May 15, 2018

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

Second-level monitoring frequency introduced since May 15, 2018

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

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

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

Restrictions

- You can configure second-level monitoring for instances that meet the following conditions:
 - The instance is located in these regions: China (Hangzhou), China (Shanghai), China (Qingdao), China (Beijing), or China (Shenzhen)

- The instance is an RDS for MySQL instance.
- The instance storage type is local SSD.
- The instance memory space is 8 GB or more.
- All engines (MySQL, SQL Server, PostgreSQL, and PPAS) and database versions support the following monitoring frequencies:
 - Once per 60 seconds
 - Once per 300 seconds

Procedure

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



Note:

Different types of databases support different metrics. For more information, see **List of monitoring items** at the end of this document.

5. Click the **Monitoring** tab.
6. Click **Set Monitoring Frequency**.
7. Select the monitoring frequency in the **Set Monitoring Frequency** dialog box and click **OK**.


Set Monitoring Frequency

Monitoring Frequency:

☐ 60 Seconds per Time ☒ 300 Seconds per Time

OK Cancel

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



Interface description:

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

List of monitoring items

RDS for MySQL

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

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

RDS for SQL Server

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

RDS for PostgreSQL

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

RDS for PPAS

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

6.2 Set monitoring rules

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

Background information

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

Procedure

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the **Basic Information** page.
4. Click **Monitoring and Alarms** in the left-side navigation pane.
5. Click the **Alarms** tab.
6. Click **Set Alarm Rules** to open the CloudMonitor console.

**Note:**

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

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

**Note:**

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

8. Click **Create Alarm Contact**.

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



Note:

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

10. On the **Alarm Contact Management** page, click the **Alarm Contact Group** tab.

11. Click **Create Alarm Contact Group**.

12. Fill in **Group Name** and **Description**, select a contact from **Existing Contacts**, click



to add the contact to **Selected Contacts**, and click **OK**.



Note:

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

13. After creating the alarm contact group, choose **Cloud Service Monitoring > ApsaraDB for RDS** from the left-side navigation pane.

14. Select the region of RDS for which the alarm rule is to be set.

15. Find the target instance and click **Alarm Rules** in the **Actions** column.

The system displays the metrics of the current alarm.

16. Click **Create Alarm Rule** to add new alarm rules.



Note:

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

7 Security

7.1 Switch the IP whitelist to enhanced security mode

IP whitelist modes

RDS instances provide two IP whitelist modes:

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

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

Changes after switching to the enhanced security mode

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

**Note:**

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

Attention

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

Procedure

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

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

| Security



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

7.2 Set SSL encryption

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




Note:

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

Enable SSL encryption

1. Log on to the [RDS Console](#).
2. Select the region where the target instance is located.
3. Click the ID of the target instance to enter the **Basic Information** page.
4. In the left-side navigation pane, click **Security** to go to the **Security** page.
5. Click the **SSL** tab.
6. Click the button next to **Disabled**, as shown in the following figure.




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




Note:

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



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



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

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

**Note:**

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

```
jdk.tls.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.



7.3 Set TDE

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



Note:

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

Background information

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

Considerations


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

Prerequisite

KMS is activated.

Procedure

1. Log on to the [RDS console](#) and select the target instance.
2. Click **Data Security** in the left-side navigation pane.
3. On the **Data Security** page, click the **TDE** tab.
4. Click **Not Activated**, as shown in the following figure.



5. Click **OK** to activate TDE.



Note:

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

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

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

Subsequent operation

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

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

8 Log management

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

Use the RDS console to manage logs

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

Procedure

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

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

Use SQL statements to manage logs

Instances in SQL Server 2012 and later versions read error logs only through the `sp_rds_read_error_logs` storage procedure. The method of using it is similar to that of using `sp_readerrorlog`.

Example 1:

```
EXEC sp_rds_read_error_logs
```

Example 2:

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

9 Backup

9.1 Back up RDS data

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

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

Backup policies

ApsaraDB supports data backup and log backup. To recover data 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. - Manual backup supports full physical backup, full logical backup, and single-database logical backup. - MySQL 5.7 Basic Edition: - Supports only snapshot-based backup instead of logical backup. - Backup files are retained for at most 7 days for free.	- After being generated, binlogs (500 MB per log) are compressed and uploaded immediately. Local files are deleted within 24 hours. - Binlog files occupy instance disk capacity. Using the binlog upload function, you can upload binlog files to OSS. This does not affect the data recovery function and stops the binlog files from occupying instance disk space.
SQL Server	- Supports full physical backup and incremental physical backup. - Automatic backup cycles from full backup, incremental backup to incremental backup. For example , if a full backup is performed on	RDS automatically generates log backups (log files). You can set the log file generation interval to 30 minutes or the data backup interval.

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

Automatic backup (Setting backup policies)

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. Click **Backup and Recovery** in the left-side navigation pane.
4. On the **Backup and Recovery** page, select **Backup Settings** and click **Edit**.

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

The parameters are explained as follows:

Back up Instance

Backup Mode : Logical Backup ▼

Backup Policy : ☐ Instance Backup ☒ Single-Database Backup

Are you sure you want to back up the instance immediately? (The backup task will start in approximately 1 minute.)

OK Cancel

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

Manual backup



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**.

Back up Instance

Backup Mode : Logical Backup

Backup Policy : ☐ Instance Backup ☒ Single-Database Backup

>

<

Are you sure you want to back up the instance immediately? (The backup task will start in approximately 1 minute.)

OK

Cancel



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**.

9.2 View the free quota of the backup space

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

Formula for calculating the free quota of the backup space

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

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

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


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

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



Note:

Instances of different types support different free quotas. The following figure is only an example.



9.3 Download data and log backup files

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

Procedure

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



Note:

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

- **Instance Number** of the binlog must be the same as **Instance Number** of the data backup.
 - The binlog backup start time must be later than the data backup time and earlier than the restoration time.
- In the **Download Instance Backup File** dialog box, select a download method.

Download Instance Backup File

Currently, backup files can be downloaded for free. In the future, the backup files will be charged. If the ECS and RDS are in the same region, using the intranet address is faster and has a higher security degree.

Methods for Downloading and Recovering Backup Files

Note: The latest version of flash is required to copy the download address.

[Download](#)

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

9.4 Logical backup and recovery for PPAS

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

Procedure

1. Install the PPAS program.

**Note:**

You must use the PPAS binary system for export. Using the PostgreSQL community binary system leads to an error.

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

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

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

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

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

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

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

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

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

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

FAQ

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

```
ERROR: permission denied for relation product_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 the PPAS binary system to export the data. For PPAS downloading methods, see the preceding procedure.

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

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

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

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

3. The following error occurs when you use pg_dump.

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

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

4. A parameter error occurs when you use pg_dump.

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

10 Recovery

10.1 Create a clone instance

To restore historical data of an instance (such as Instance A), you can restore the data to a clone instance, verify the data on the clone instance, and transfer the data you need from the clone instance to the master instance (Instance A). This article describes how to restore data to a clone instance by creating a clone instance. The restored data includes instance data and settings.

Clone instances are managed and billed in the same way as the master instance. For information about the cost of instances, see [Pricing](#).

**Note:**

Currently, the following RDS versions support clone instances:

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

Background information

You can specify a backup set or any time point within the backup period to create a clone instance. Clone instances only copy the content of the master instance. They do not copy the content of read-only instances or disaster recovery instances under the master instance. The copied content includes information about databases, accounts, and instance settings (such as whitelists, backup settings, parameters, and alarm thresholds).

The database engine of a clone instance must be the same as that of the master instance. Other settings can be different, such as the billing method, instance series, zone, network type, instance specifications, and storage capacity. If a clone instance is used to recover data of the master instance, we recommend that the clone instance is configured with higher specifications and storage capacity than the master instance. Otherwise, the recovery may take a long time.

The accounts of the clone instance are the same as those of the master instance, but you can modify the account passwords. For example, if you create a clone instance for a master instance that uses a master account, the clone instance also uses the master account.

Prerequisites

The master instance must meet the following conditions:

- The instance status is running and unlocked.
- No migration task is ongoing.
- Data backup and log backup are enabled.
- If the clone instance is to be created from a backup set, the master instance must have at least one backup set.

**Note:**

To use a sub-account to create a clone instance, ensure that the sub-account has added authorization policies for the clone instance. For details about how to authorize, see [Authorization for RDS 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 go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Recovery**.
5. At the upper-right corner of the page, click **Restore Database**.
6. Select a payment method: **Subscription** or **Pay-As-You-Go**.
7. Configure the clone instance.

Clone Instance

Restore Type :

By TimeBy Backup ID

Backup ID :

Select a backup set

Edition :

Zone :

China East 1 Zone B

Type :

1 Core, 1 GB

(Instance Type Code: rds.mysql.t

The maximum number of connections :300 ; IOPS : 600

Storage

Capacity :

500GB

1000GB

2000GB

Network

Network Type :

Classic NetworkVPC

Learn more>> ?

Purchase Plan

Duration :

1 month

2

3

4

5

6

7

8

9

1Yr

Quantity :

1

Parameter descriptions:

Parameter Name	Description
Restore Type	Restore data by time or by backup set.

Parameter Name	Description
Restored At	This parameter is available if Restore Type is set to By Time . You can set the parameter to any point in time.
Backup ID	This parameter is available if Restore Type is set to By Backup ID . Select a backup set.
Edition, Zone, Type, Capacity, Network Type, and Duration	For descriptions of these parameters, see Create an instance .
Quantity	You can create up to five clone instances in one order.

8. Click **Buy Now**.
9. Select **Product Terms of Service and Service Level Notice and Terms of Use**, and click **Pay Now** to confirm the order information.

10.2 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, 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 method is only applicable to databases 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. Click **Backup and Recovery** in the left-side menu.
5. Click the **Backup List** tab.
6. Select a time range for recovery and click **Query**.
7. Select the target backup file and click **Coverage Restoration**, as shown in the following figure.

Backup List	Binlog List	Temporary Instance	Backup Settings	
Select Time Range	2017-05-09	To	2017-05-16	Query
Backup Start/End Time	Backup Policy	Backup Size		
2017-05-16 15:39/2017-05-16 15:41	Instance Backup	1 MB		

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

11 Create a linked server for SQL Server instances

This document is applicable only to high-availability instances of RDS for 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 storage procedures is complex.
- You cannot create a linked server using DNS and the corresponding IP address.

Despite the constraints, this document provides a simple method of 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 is displayed 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.

12 Data migration

12.1 Migrate from other cloud databases to ApsaraDB RDS

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

Migrate MySQL from AWS RDS to ApsaraDB RDS

[Migrate MySQL from AWS RDS to ApsaraDB RDS with DTS](#)

12.2 Use mysqldump to migrate MySQL data

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

Background information

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

Prerequisites

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

Procedure

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

1. Create a migration account in the local database.

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

Parameter description:

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

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

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

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

```
GRANT SELECT ON 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 host, you can use the wildcard *%*.
- **WITH GRANT OPTION**: indicates an optional parameter that enables the account to use the *GRANT* command.

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

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

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



Note:

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

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

Parameter description:

- **localIp**: IP address of the local database server

- **userName**: Migration account of the local database
- **dbName**: Name of the database to be migrated
- **/tmp/dbName.sql**: Backup file name

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



Note:

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

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

Parameter description:

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

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

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

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

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

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

Parameter description:

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

12.3 Migrate RDS data to the local database

12.3.1 Migrate RDS for MySQL data to a local MySQL database

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

Export based on a physical backup file

Background information

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

Prerequisites

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

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

For installation instructions, see [Percona XtraBackup 2.3](#) and [Percona XtraBackup 2.4](#).

Procedure

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

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

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

3. Decompress the backup file.

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

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

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

```
cd filename/
```

```
ll
```

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

```
-rw-r--r-- 1 root root      269 Aug 19 18:15 backup-my.cnf
drwxr-xr-x 2 root root    4096 Aug 21 10:31 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 based on a logical backup file

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

Procedure

1. Download **logical backup file** and upload the file to the target server. For more information about how to obtain the backup file, see [Download RDS data and log backup](#). If the target server can access the source instance, you can use `wget "url"` to download the backup file. `url` indicates the backup file downloading address.

2. Switch to the backup file path.

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

3. Decompress the backup file.

```
tar xzvf 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.

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

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

Procedure

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

For more information about how to obtain the backup file, see [Download RDS data and log backup](#).

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

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

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

- datafull: indicates full backup, such as **rdsumu2myfzbeai1_datafull_201402250050_2250050.bak**.
- datadiff: indicates incremental backup, such as **rdsumu2myfzbeai1_datadiff_201402260050_2260050.bak**.
- log: indicates 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 for saving the full backup file: *d:\backup\rdsumu2myfzbeai1_datafull_201402250050_2250050.bak*

- Path for saving 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='d:\backup\rdsumu2myfzbeai1_datafull_201402250050_2250050.bak';
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**.

SQLQuery1.sql - (Administrator (52))*

```
restore filelistonly from disk='D:\backup\rdsumu2myfzbeai1_datafull_201402250050_2250050.bak';
go
```

	LogicalName	PhysicalName	Type	FileGroupName	Size
1	data1	d:\MS3004\Data\rdsumu2myfzbeai1\data1.mdf	D	PRIMARY	53
2	log	d:\MS3004\Log\rdsumu2myfzbeai1\log.ldf	L	NULL	53

5. Load the full backup file.

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

Parameters description:

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

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

**Note:**

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

6. Load the incremental backup file.

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

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

7. Recover the database.

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

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

12.3.3 Migrate RDS for PostgreSQL data to a local PostgreSQL database

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

Procedure

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

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

Parameters description:

- **username**: indicates the user name used for database logon.
- **hostname**: indicates the host name of the database.
- **port**: indicates the database port number.
- **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 dumpfilename.sql
```

Parameter description:

- **username**: indicates the user name 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 data import. They can be ignored, for example:

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

12.3.4 Migrate RDS for PPAS data to a local Oracle database

Constraint

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

Prerequisites

- An Oracle database must be installed on the server.
- The IP address of the Oracle server must be added to the whitelist of the RDS for PPAS database instance. For specific instructions, see [Set whitelist](#).
- You must create a table structure in Oracle that corresponds to the RDS for PPAS database table structure.

- The PostgreSQL client has been uploaded to the Oracle database server.

Procedure



Note:

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

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

```
[root@oraclexe ~]# yum install postgresql.x86_64
[root@oraclexe ~]# /usr/bin/psql --version
psql (PostgreSQL) 8.4.20
```

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

```
[root@oraclexe ~]# vim ~/.pgpass
[root@oraclexe ~]# cat ~/.pgpass
rm-2ze46615u1k657yyn.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-2ze46615u1k657yyn.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 user **ora**, the connection has been established. After a successful test, return to user **root**.

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

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

- a. Create a file `ppas_exp_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 the Oracle database.

- Method 1: Use Oracle SQL Loader to import data. For more information, see [Oracle SQL Loader Overview](#).
- Method 2: Use Oracle SQL Developer to import data. For more information, see [SQL Developer Concepts and Usage](#).

Troubleshooting

Problem

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

```
[root@oraclexe ~]# ./ppas_exp_all_tables_to_csv.sh rm-2ze46615u1
k657yyn.ppas.rds.aliyuncs.com 3433 myadmin ora
mkdir: Cannot create directory: "/tmp/ppas_tables_rm-2ze46615u1k657yyn
.ppas.rds.aliyuncs.com_3433_myadmin_ora": file already exists
```

Solution

Delete the existing directory.

```
[root@oraclexe ~]# rm -rf /tmp/ppas_tables_rm-2ze46615u1k657yyn.ppas.
rds.aliyuncs.com_3433_myadmin_ora
```

12.3.5 Migrate RDS for PPAS data to a local PPAS database

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

Procedure

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

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

Parameter descriptions:

- **username**: indicates the user name used for database logon.
- **hostname**: indicates the host name of the database.
- **port**: indicates the database port number.

- **dbname**: 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 rds-v07z563m7o25cj550public.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 destinationdb -p port -f dumpfilename.sql
```

Parameter descriptions:

- **username**: indicates the user name used for database login.
- **hostname**: indicates the database address.
- **port**: indicates the database port number.
- **dbname**: indicates the database name.
- **filename**: indicates the backup file name. For example:

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

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

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

12.4 Compress data

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

Restrictions

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

Procedure

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

```
SELECT version();
```



Note:

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

2. Set the proportion of **loose_tokudb_buffer_pool_ratio**, namely, 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.

12.5 Use psql to migrate PostgreSQL data

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

Background information

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

Prerequisite

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

Prepare local data

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

```
pg_dump -U username -h hostname -p port 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
```

Migrate data



Note:

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

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

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

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

12.6 Migrate SQL Server to cloud

12.6.1 Migrate data to RDS for SQL Server 2008 R2

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

Prerequisite

A target database has been created in RDS. For more information, see [Create database and account for SQL Server 2008 R2](#).



Note:

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

Billing details

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




Figure description:

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



Note:

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

Procedure

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

```
use master;
go
```

```
select name, case recovery_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, go to Step d.
- If the model value is FULL, go to Step e.

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



Note:

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

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

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

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

f. Run the following commands to verify 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. In this case, back up the database again.

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



Note:

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


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

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

a. Upload the backup file to OSS:


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

b. In the left-side navigation pane of the [OSS console](#), select the bucket where the backup file belongs.



c. Select **Files**.

d. Click the name of the target backup file.



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



Note:

When you migrate the backup file from OSS to RDS, the URL of the backup file is required. If the link validity period for the URL expires, the data migration fails. Therefore, we recommend that you set the validity period to the maximum value, which is 28,800s.

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

The screenshot shows a configuration form with the following elements:

- File Name:** A text input field with a blurred value.
- Signature:** A section containing a help icon (?) and a sub-field labeled "Validity (seconds)" with the value "28800".
- URL:** A text input field containing the URL: `http://cjltest-intl.oss-cn-hangzhou.aliyuncs.com/Penguins.jpg?Expires=1515585060&OSSAccessKeyId=TMP.AQHFsKphKVozDyCHYy5YrSzcfD9X8JEI3vOc`.
- Buttons:** Two buttons are located below the URL field: "Copy File URL" (which is highlighted with a red rectangular box) and "Copy File Path".

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

For example, if the backup file URL is `http://rdstest-yanhua.oss-cn-shanghai.aliyuncs.com/testmigraterds_20170906143807_FULL.bak?Expires=1514189963&OSSAccessKeyId=TMP.AQGVf994YTPfArSpw78uix2rdGBi-dPe_FzQSLw0LP7MVlR-XXXX`, change the Internet endpoint `oss-cn-shanghai.aliyuncs.com` in the URL to the intranet endpoint `oss-cn-shanghai-internal.aliyuncs.com`.

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

Database Name	Database Status	Character Set	Bound Accounts	Description	Action
test	Running	Chinese_PRC_CI_AS		None	Delete Migrate backup files from OSS

- 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**.

Import Guide
✕

1. Back up your database
2. Upload the backup files
3. Import data

Database Name

OSS URL of the Backup File

Cloud Migration Plan
☒ One-time full backup file migration

Exit the Wizard
Previous
OK

- i. Click **OK**.
- j. In the left-side navigation pane, click **Data Migration to Cloud** to go to the page listing the tasks of migrating backup files from OSS to RDS.
- k. Find the target migration task. If **Tasks Status** is **Success**, the data is successfully migrated to the RDS database. If the migration task status does not change to **Success** after a long time, click **View File Details** next to the migration task to view the failure

causes. After solving the problems, perform the required steps to migrate the backup file again.

13 Typical applications

13.1 Cached data persistence

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

Background information

Compared with RDS, ApsaraDB for Memcache and the ApsaraDB for Redis have the following two features:

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

System requirements

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

bmemcached download address: Click [Here](#) to download.

The bmemcached installation command is as follows:

```
pip install python-binary-memcached
```

- Python is used as an example. Python and pip must be installed in the local environment or ECS.

Sample code

The following sample code realizes the combined use of 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
```

13.2 Multi-structure data storage

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

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

Sample code

OSS works with the RDS.

1. Initialize OssAPI.

```
from oss.oss_api import *
endpoint="oss-cn-hangzhou.aliyuncs.com"
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')
```