

# Alibaba Cloud ApsaraDB for MySQL

## Product Introduction

Issue: 20180806

# 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.	 <b>Danger:</b> 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.	 <b>Warning:</b> 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.	 <b>Note:</b> 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.	 <b>Note:</b> You can use <b>Ctrl + A</b> to select all files.
>	Multi-level menu cascade.	<b>Settings &gt; Network &gt; Set network type</b>
<b>Bold</b>	It is used for buttons, menus, page names, and other UI elements.	Click <b>OK</b> .
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

---

<b>Legal disclaimer.....</b>	<b>I</b>
<b>Generic conventions.....</b>	<b>I</b>
<b>1 What is RDS.....</b>	<b>1</b>
<b>2 Benefits.....</b>	<b>3</b>
2.1 Inexpensive and easy-to-use.....	3
2.2 High performance.....	3
2.3 High security.....	4
2.4 High reliability.....	5
2.5 Comparison between ApsaraDB for RDS and local databases.....	6
<b>3 System architecture.....</b>	<b>9</b>
3.1 Data link service.....	9
3.2 High-availability service.....	10
3.3 Backup and recovery service.....	18
3.4 Monitoring service.....	18
3.5 Scheduling service.....	19
3.6 Migration service.....	20
<b>4 Product series.....</b>	<b>21</b>
4.1 General introduction to product series.....	21
4.2 Basic Edition.....	22
<b>5 Instance types.....</b>	<b>25</b>
5.1 Instance type overview.....	25
5.2 Dedicated instance.....	26
5.3 Instance type list.....	27
<b>6 Typical applications.....</b>	<b>39</b>
6.1 Diversified data storage.....	39
6.2 Read/write splitting function.....	40
<b>7 RDS usage instructions.....</b>	<b>42</b>

# 1 What is RDS

---

ApsaraDB for RDS (Relational Database Service) is a stable and reliable online database service, and it also supports elastic scaling function. RDS is a stable, reliable and flexible online database service. Based on the Apsara distributed system and high-performance storage of ephemeral SSD, RDS supports MySQL, SQL Server, and PostgreSQL engines. It offers a complete set of solutions for backup, recovery, monitoring, migration, disaster recovery, and troubleshooting database operation and maintenance.

## ApsaraDB for MySQL

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

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

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

RDS for MySQL currently supports versions 5.5, 5.6, and 5.7. For specific support details, see the Quick Start guide of your corresponding MySQL version.

## ApsaraDB for SQL Server

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

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

RDS for SQL Server currently supports the following versions:

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

### **ApsaraDB for PostgreSQL**

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

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

RDS for PostgreSQL currently supports version 9.4.

### **ApsaraDB for PPAS**

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

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

RDS for PPAS currently supports version 9.3.

## 2 Benefits

---

### 2.1 Inexpensive and easy-to-use

#### Simple deployment

You can customize RDS specifications through Alibaba Cloud's official website or the API. After the order is confirmed, RDS generates the specified instance instantly. RDS can work with ECS to reduce the application response time and save on public traffic fees.

#### On-demand upgrades

Initially, you can purchase an RDS instance that meets the existing business requirements. When requirements on the database and data storage capacity change, you can flexibly adjust the instance specifications without any interruptions to the service.

#### Effortless migration

RDS is used similarly to the native database engine, meaning that you can transfer the pre-existing knowledge and skills to RDS management. Data can be migrated to RDS using the commercial off-the-shelf data import and export tools with minimal labor cost required.

#### Ease of management

Alibaba Cloud is responsible for ensuring the normal operation of RDS through routine maintenance and management, such as hardware/software fault processing and database update patches. You can independently perform database addition, deletion, restart, backup, recovery, and other management operations in the Alibaba Cloud console.

#### Related Topics

- [High performance](#)
- [High Security](#)
- [High reliability](#)
- [Comparison between ApsaraDB for RDS and local databases](#)

### 2.2 High performance

#### Parameter optimization

Alibaba Cloud has accumulated years of experience in production and optimization by gathering key opinions from top database experts in China and aggregating performance data on all the

RDS instances. DBA continuously manages RDS over its lifecycle to ensure that RDS is running at the optimal performance.

### SQL optimization

Based on your application scenario, RDS locks low-efficiency SQL statements and offers recommendations for optimizing your business code.

### High-end backend hardware

All servers used by RDS go through multiple levels of service verification by multiple parties to ensure the exceptional performance and stability.

### Extended reading

- [Instance type list](#)
- [Performance White Paper \(RDS for MySQL\)](#)
- [Performance White Paper \(RDS for SQL Server\)](#)
- [Precautions on comparing ECS user-created databases and ApsaraDB for RDS instances in terms of performance](#)

## 2.3 High security

### Anti-DDoS attack

When Internet connection is used to access RDS instances, the risk of DDoS attacks occurring on the network is possible. If this occurs, the RDS security system enables flow cleaning operation first. If the flow cleaning operation fails or the attack reaches the black hole threshold, black hole processing is triggered. Please refer to attack [Protection and attack protection](#) for details.

**Notice:** We recommend that RDS instances are accessed over the intranet to avoid DDoS attacks.

### Access control policy

- You can define the IP addresses that are allowed to access RDS. IP addresses that have not been specified are denied access.
- Each account can only view and operate its own database.

Please refer to [Access control](#) for more information.

## System security

- RDS is protected by multiple firewall layers that can effectively block a variety of malicious attacks and guarantee data security.
- Direct logon to the RDS server is not allowed. Only the port required by the specific database service is open.
- The RDS server cannot initiate an external connection. It can only accept access requests.

For more information, refer to [Access control](#).

## Professional support team

Alibaba Cloud's security team provide rapid security technology support for RDS.

### Related Topics

- [Cheap and ease-to-use](#)
- [High performance](#)
- [High reliability](#)
- [Comparison between ApsaraDB for RDS and local databases](#)

## 2.4 High reliability

### Hot standby

RDS uses hot standby, so if a physical server fails, the service is switched over in seconds without interruptions to application services.

### Multi-copy redundancy

The data in RDS server is built on RAID, and the data backup is stored on OSS.

### Data backup

RDS offers an automatic backup mechanism. You can set a backup schedule or initiate a temporary backup at any time. For more information, see [Backup recovery](#).

### Data recovery

Data can be recovered from a backup. Generally, data can be recovered within 7 days to a temporary RDS instance. After the data is verified, the data can be migrated back to the master RDS instance. For more information, see [Backup recovery](#).

## 2.5 Comparison between ApsaraDB for RDS and local databases

### Performance comparison

Item	ApsaraDB for RDS	Database services built on purchased servers
Service availability	99.95%	You are required to implement your own data protection, master-slave replication, and RAID.
Data reliability	99.9999%	You are required to implement your own data protection, master-slave replication, and RAID.
System security	Anti-DDoS, malicious traffic cleaning, timely repair of various database security vulnerabilities	You are required to do deployment at high costs and repair database security vulnerabilities independently.
Database backup	Automatic backup	Independent backups are possible, but you have to find storage space for backup and regularly validate whether the backup data can be recovered.
Software and hardware investment	Pay as you go, with no software and hardware investment	Database servers are rather costly and the license fees for SQL Server have to be paid.
System hosting	No hosting fee	A single 2U server costs more than RMB 5,000 RMB/year (more than 10,000 RMB/year for master and slave servers, if needed).
Maintenance cost	Maintenance not needed	Professional DBAs have to be hired, resulting in high labor cost.
Deployment and resizing	Instant activation, fast deployment, auto resizing, activation on demand	Hardware procurement, data center hosting, and host deployment, among other tasks, are required and can be time consuming.
Resource utilization rate	Charged based on actual usage, resulting in 100% utilization.	Resource utilization is low if peak traffic is considered.

## Price comparison

Item	ApsaraDB for RDS	Database services built on purchased servers
Costs of hardware, spare parts and accessories	Take the following instance type as an example: An instance with memory of 1,200 MB and storage space of 50 GB (IOPS is up to 600) costs 2,040 RMB/year.	<ul style="list-style-type: none"> <li>A minimum of 2 servers are required for the database cluster, and a single server whose IOPS is up to 600 costs approximately 6,000 RMB.</li> <li>An intranet switch is used to connect the front-end Web server (an inexpensive 1U non-NMS switch costs about 1,000 RMB). Later hardware damages and replacements take at least 30% of the cost.</li> <li>Hardware cost: <math>(6,000 \times 2 + 1,000) \times 130\% = 16,900</math> RMB</li> </ul> <p>Annual hardware cost: 16,900 RMB/3= 5,633 RMB (the cost of hardware is calculated over a 3-year depreciation period)</p>
Data center hosting fee	It is service provider's responsibilities, and no hosting fee is required.	The hosting fee for a 1U cabinet is 3,000 RMB/year, and the hosting fee for two 1U servers and a 1U intranet switch is charged. Data center hosting fees: $3,000 \times 3 = 9,000$ RMB
Bandwidth fee	Communication between ECS and ApsaraDB for RDS in the same region is available through the intranet at no cost. However, communication between ECS and ApsaraDB for RDS across different regions is available over the Internet at a certain cost for Internet traffic. For more information, see <a href="#">Pricing</a> .	Available only in the intranet at no cost for Internet traffic.
Cost of database maintenance engineers	There is no labor cost, because database maintenance is taken care of by a service provider.	The monthly salary for an entry-level DBA engineer is at least 5,000 RMB/month. If 30% workload for a full-time

Item	ApsaraDB for RDS	Database services built on purchased servers
		engineer is needed for the current project: Labor cost: $5,000 \times 12 \times 30\%$ = 18,000 RMB
Total annual cost	2,040 RMB/year	32,633 RMB/year

## 3 System architecture

---

### 3.1 Data link service

ApsaraDB provides all of the data link services, including DNS, SLB, and Proxy. Because RDS uses the NativeDB Engine, and database operations are highly similar across engines, no learning cost is generated for users who are familiar with these data link services.

#### DNS

The DNS module supports the dynamic resolution of domain names to IP addresses, to prevent IP address changes from affecting the performance of RDS instances. After its domain name is configured in the connection pool, an ApsaraDB instance continues to be accessed even if the corresponding IP address changes.

For example, the domain name of an ApsaraDB instance is `test.rds.aliyun.com`, and the IP address corresponding to this domain name is `10.10.10.1`. If either `test.rds.aliyun.com` or `10.10.10.1` is configured in the connection pool of a program, the instance can be accessed.

After performing a zone migration or version upgrade for this ApsaraDB instance, the IP address may change to `10.10.10.2`. If the domain name configured in the connection pool is `test.rds.aliyun.com`, the instance can still be accessed. However, if the IP address configured in the connection pool is `10.10.10.1`, the instance is no longer accessible.

#### SLB

The SLB module provides instance IP addresses (including both intranet and Internet IP addresses) to prevent physical server changes from affecting the performance of RDS instances.

For example, the intranet IP address of an RDS instance is `10.1.1.1`, and the corresponding Proxy or DB Engine runs on `192.168.0.1`. Normally, the SLB module redirects all traffic destined for `10.1.1.1` to `192.168.0.1`. If `192.168.0.1` fails, another address in hot standby status, `192.168.0.2`, takes over for `192.168.0.1`. In this case, the SLB module redirects all traffic destined for `10.1.1.1` to `192.168.0.2`, and the RDS instance continues to offer its services normally.

#### Proxy

The Proxy module performs a number of functions including data routing, traffic detection, and session holding.

- **Data routing:** This supports distributed complex query aggregation for big data and provides the corresponding capacity management.

- Traffic detection: This reduces SQL injection risks and supports SQL log backtracking when necessary.
- Session holding: This prevents database connection interruptions if any failure occurs.

## DB engine

RDS fully supports mainstream database protocols, as shown in the following table:

Database type	Version
MySQL	5.1 (Deprecated), 5.5, 5.6, 5.7
SQL Server	2008 R2, 2012, 2016
PostgreSQL	9.4
PPAS	9.3, highly compatible with Oracle

## 3.2 High-availability service

The high-availability service consists of several modules including the Detection, Repair, and Notification modules. In combination, these modules guarantee the availability of the data link services and process any internal database exception.

In addition, RDS can improve the performance of its high-availability service by migrating to a region that supports multiple zones and by adopting the appropriate high-availability policies.

### Detection

The Detection module checks whether the master and slave nodes of the DB Engine offer their services normally. The HA (High Available) node uses heartbeat information, acquired at an interval of 8 to 10 seconds, to check the health status of the master node. This information, combined with the health status of the slave node and heartbeat information from other HA nodes, allows the Detection module to eliminate any risk of misjudgment caused by exceptions such as network jitter and allows that the exception switchover can be completed within 30 seconds.

### Repair

The Repair module maintains the replication relationship between the master and slave nodes of the DB Engine. It can also repair any errors that may occur on either node.

For example:

- Automatic restoration of master/slave replication in case of disconnection
- Automatic repair of table-level damage to the master or slave nodes

- On-site saving and automatic repair if the master or slave nodes crash

## Notice

The Notice module informs the SLB or Proxy of status changes to the master and slave nodes to guarantee that you can continue to access the correct node.

For example, the Detection module discovers that the master node has an exception and instructs the Repair module to fix it. If the Repair module fails to resolve the problem, it directs the Notification module to initiate traffic switching. The Notification module then forwards the switching request to the SLB or Proxy, which begins to redirect all traffic to the slave node. Simultaneously, the Repair module creates a new slave node on another physical server and synchronizes this change back to the Detection module. The Detection module then incorporates this new information and starts to recheck the health status of the instance.

## Multi-zone

Multi-zone refers to the physical area that is formed by combining multiple individual zones within the same region. Multi-zone RDS instances can withstand higher level disasters than single-zone instances. For example, a single-zone RDS instance can withstand server and rack failures, while a multi-zone RDS instance can survive a situation such as failure of an entire data center.

Currently no extra charge for multi-zone RDS instances is generated. Users in a region where multi-zone is enabled can purchase multi-zone RDS instances directly or convert single-zone RDS instances into multi-zone RDS instances by using inter-zone migration.



### Note:

Multiple zones may have a certain amount of network latency. As a result, when a multi-zone RDS instance uses a semi-synchronous data replication solution, its response time to any individual update may be longer than that of a single-zone instance. In this case, the best way to improve overall throughput is to increase concurrency.

## High-availability policy

The high-availability policies use a combination of service priorities and data replication modes to meet the business needs.

The service priorities are as follows:

- RTO (Recovery Time Objective) priority: The database must restore services as soon as possible within a specified time frame. This is best for users who require their databases to provide uninterrupted online service.
- RPO (Recovery Point Objective) priority: The database must guarantee the data reliability, that is, as little data loss as possible. This is best for users whose highest priority is data consistency.

There are three data replication methods:

- Asynchronous replication (Async):

In this mode, the master node does not immediately synchronize data to the slave node. When an application initiates an update request, which may include add, delete, or modify operations, the master node responds to the application immediately after completing the operation but does not necessarily replicate that data to the slave node right away. This means that the operation of the primary database is not affected if the slave node is unavailable, but data inconsistencies may occur if the master node is unavailable.

- Forced synchronous replication (Sync):

When an application initiates an update (add, delete, or modify) request, the primary database completes the operation and then replicates data to the standby database. After the standby database receives the data, it returns a success message to the primary database. The primary database waits for a feedback from the standby database before responding to the application. This means that the operation of the master node is affected if the slave node is unavailable, but the data on the master and slave nodes is always consistent.

- Semi-synchronous replication (Semi-Sync):

This functions as a hybrid of the two preceding replication modes. In this mode, when both nodes are functioning normally, data replication is identical to the forced synchronous replication mode. However, when there is an exception, such as the slave node becoming unavailable or a network exception occurring between the two nodes, the master node only attempts to replicate data to the slave node and suspend its response to the application for a set period of time. Once the replication mode has timed out, the master node degrades to asynchronous replication. At this point, if the master node becomes unavailable and the application updates its data from the slave node, it is consistent with the data on the master node. When data replication between the two nodes returns to normal, because the slave node or network connection is recovered, forced synchronous replication is reinstated. The amount of time it takes for the nodes to return to forced synchronous replication depends on how the semi-

synchronous replication mode was implemented. For instance, ApsaraDB for MySQL 5.5 is different from ApsaraDB for MySQL 5.6 in this regard.

Several combinations of service priorities and data replication modes are available to meet your database and business needs. The characteristics of key combinations are detailed in the following table.

Cloud data engine	Service priority	Data replication mode	Combination characteristics
MySQL 5.1	RPO	Async	<ul style="list-style-type: none"> <li>If the master node fails, the slave node switches over after applying all of the relay logs.</li> <li>If the slave node fails, application operations on the master node are not affected. The data on the master node is synchronized after the slave node recovers.</li> </ul>
MySQL 5.5	RPO	Async	<ul style="list-style-type: none"> <li>If the master node fails, the slave node switches over after applying all of the relay logs.</li> <li>If the slave node fails, application operations on the master node are not affected. The data on the master node is synchronized after the slave node recovers.</li> </ul>
MySQL 5.5	RTO	Semi-Sync	<ul style="list-style-type: none"> <li>If the master node fails and data replication degrades,</li> </ul>

Cloud data engine	Service priority	Data replication mode	Combination characteristics
			<p>RDS immediately triggers the switchover and direct traffic to the slave node because data consistency is guaranteed.</p> <ul style="list-style-type: none"> <li>If the slave node fails, application operations on the master node times out, and data replication degrades to asynchronous replication. After the slave node recovers and the data on the master node is synchronized completely, data replication returns to forced synchronization.</li> <li>If the master node fails while the two nodes have inconsistent data and the data replication mode degrades to asynchronous replication, the slave node switches over after applying all of the relay logs.</li> </ul>
MySQL 5.6	RPO	Async	<ul style="list-style-type: none"> <li>If the master node fails, the slave node</li> </ul>

Cloud data engine	Service priority	Data replication mode	Combination characteristics
			<p>switches over after applying all of the relay logs.</p> <ul style="list-style-type: none"> <li>If the slave node fails, application operations on the master node are not affected. The data on the master node is synchronized after the slave node recovers.</li> </ul>
MySQL 5.6	RTO	Semi-Sync	<ul style="list-style-type: none"> <li>If the master node fails and data replication degrades, RDS immediately triggers the switchover and direct traffic to the slave node because data consistency is guaranteed.</li> <li>If the slave node fails, application operations on the master node times out, and data replication degrades to asynchronous replication. After the slave node recovers and the data on the master node is synchronized completely, data replication returns to forced synchronization.</li> </ul>

Cloud data engine	Service priority	Data replication mode	Combination characteristics
			<ul style="list-style-type: none"> <li>If the master node fails while the two nodes have inconsistent data and the data replication mode degrades to asynchronous replication, the slave node switches over after applying all of the relay logs.</li> </ul>
MySQL 5.6	RPO	Semi-Sync	<ul style="list-style-type: none"> <li>If the master node fails and data replication has not degraded, RDS immediately triggers the switchover and direct traffic to the slave node because data consistency has been guaranteed.</li> <li>If the slave node fails, application operations on the master node times out, and data replication degrades to asynchronous replication. When the slave node can obtain information from the master node again, because the slave node or</li> </ul>

Cloud data engine	Service priority	Data replication mode	Combination characteristics
			<p>network connection recovers, data replication returns to forced synchronization.</p> <ul style="list-style-type: none"> <li>If the master node fails while the two nodes have inconsistent data and the data difference on the slave node cannot be reconciled completely, you can obtain the time of the slave node through the API. Then you can decide when to switchover and which method you plan to use to reconcile the data.</li> </ul>
MySQL 5.7	X	X	Currently this engine does not support policy adjustments.
SQL Server 2008 R2	X	X	Currently this engine does not support policy adjustments.
SQL Server 2012	X	X	Currently this engine does not support policy adjustments.
PostgreSQL	X	X	Currently this engine does not support policy adjustments.
PPAS	X	X	Currently this engine does not support policy adjustments.

## 3.3 Backup and recovery service

This service supports the offline backup, dumping, and recovery of data.

### Backup

The Backup module compresses and uploads the data and logs on both the master and slave nodes.

RDS uploads this data to OSS by default, but the backup files can also be dumped to a cheaper and more persistent Archive Storage. Normally, backup is initiated on the slave node so as not to affect the performance of the master node. However, if the slave node is unavailable or damaged, the Backup module initiates backup on the master node.

### Recovery

The Recovery module recovers a backup file stored on OSS to the target node.

- Master node roll-back: This can be used to restore a node to the state that it was in at a specific point in time.
- Slave node repair: This can be used to automatically create a new slave node to reduce risks when an irreparable failure occurs to the slave node.
- Read-only instance creation: This creates a read-only instance from a backup.

### Storage

The Storage module is responsible for uploading, dumping, and downloading backup files.

Currently all backup data is uploaded to OSS for storage, and users can obtain temporary links to download their data as needed. The Storage module also supports dumping backup files from OSS to an Archive Storage for cheaper and steady offline storage.

## 3.4 Monitoring service

ApsaraDB provides multilevel monitoring services across the physical, network, and application layers to ensure business availability.

### Service

The Service module tracks service-level statuses. It monitors whether the SLB, OSS, Archive Storage, Log Service, and other cloud products on which RDS depends are normal, including their functionality and response time. It also uses logs to determine whether the internal RDS services are operating normally.

## Network

The Network module tracks statuses at the network layer. It monitors the connectivity between ECS and RDS and between physical RDS servers, as well as the rates of packet loss on the router and switch.

## OS

The OS module tracks statuses at the hardware and OS kernel layer, including:

- Hardware overhaul: Constantly checks the operational status of the CPU, memory, main board , and storage, pre-judges whether a fault will occur, and automatically submits a repair report in advance.
- Hardware overhaul: Constantly checks the operational status of the CPU, memory, main board , and storage, pre-judges whether a fault will occur, and automatically submits a repair report in advance.

## Instance

The Instance module collects RDS instance-level information, including:

- Available instance information.
- Instance capacity and performance indicators.
- Instance SQL execution records.

# 3.5 Scheduling service

The scheduling service consists of a Resource module and a Version module. It mainly implements resource allocation and instance version management.

## Resource

The Resource module allocates and integrates the underlying RDS resources. For example, when you create an instance through the [RDS console](#) or API, the Resource module determines which physical server is best suited to carry traffic. This module also allocates and integrates the underlying resources required for the inter-zone migration of RDS instances. After lengthy instance creation, deletion, and migration operations, the Resource module calculates the degree of resource fragmentation in a zone and initiates resource integration regularly to improve the service carrying capacity of the zone.

## Version

The Version module is applicable to version upgrades of RDS instances. For example:

- MySQL major version upgrade: Upgrade MySQL 5.1 to MySQL 5.5, and upgrade MySQL 5.5 to MySQL 5.6.
- MySQL minor version upgrade: Fix a bug in the MySQL source code.

## 3.6 Migration service

The migration service can migrate data from a local database to ApsaraDB, or migrate an ApsaraDB instance to another instance. ApsaraDB includes a Data Transfer Service (DTS) tool to facilitate quick database migration.

DTS is a cloud data transfer service for efficient instance migration in local databases, and between RDS instances. For more information, see [DTS product overview](#).

DTS provides three migration modes, that is, structural migration, full migration, and incremental migration. The detailed introduction is as follows:

- Structural migration: DTS migrates the structural definitions of migration objects to the target instance. Currently, this mode can migrate tables, views, triggers, stored procedures, and stored functions.
- Full migration: DTS migrates all existing data in the source database migration objects to the target instance.

**Note:**

To guarantee data consistency, non-transaction tables that do not have a Primary Key are locked during full migration. Locked tables cannot be written to, and the duration of the lock depends on the volume of data in a table. The locks are released only after the non-transaction tables without a Primary Key are fully migrated.

- Incremental migration: DTS synchronizes data changes made in the migration process to the target instance.

**Note:**

If a DDL operation is performed during data migration, structural changes are not synchronized to the target instance.

## 4 Product series

### 4.1 General introduction to product series

Currently, RDS instances are divided into three series: Basic Edition, High-availability Edition and Finance Edition. Different series support different engine types and instance types. See [Instance type overview](#) for more information.

**Note:**

Currently the Finance Edition is applicable only to the regions in China.

#### Brief introductions

Series	Introduction	Use case
Basic Edition	It uses the storage-computing isolated architecture and a single computing node, which realizes super high cost effectiveness.	<ul style="list-style-type: none"><li>• Personal learning.</li><li>• Micro websites.</li><li>• Development and testing environments for and small and medium-sized enterprises.</li></ul>
High-availability Edition	It uses the classic high-availability architecture with one master node and one slave node. The ephemeral SSD storage provides the best and balanced performance.	<ul style="list-style-type: none"><li>• Production databases of large enterprises.</li><li>• Applications for the industries of Internet, Internet of things (IoT), retail e-commerce sales, logistics, game, and other industries.</li></ul>

#### Feature differences

Feature	Basic Edition	High-availability Edition
Engine	<ul style="list-style-type: none"><li>• MySQL 5.7</li><li>• SQL Server 2012 Web, Enterprise</li><li>• SQL Server 2016 Web</li></ul>	<ul style="list-style-type: none"><li>• MySQL 5.5/5.6/5.7</li><li>• SQL Server 2008 R2 Enterprise</li><li>• SQL Server 2012 Standard , Enterprise</li><li>• SQL Server 2016 Standard , Enterprise</li></ul>

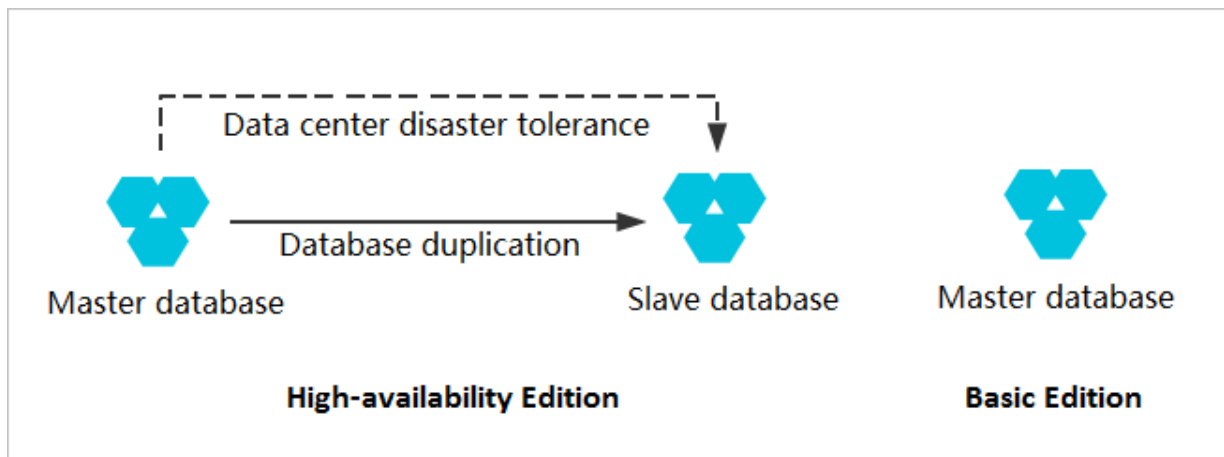
Feature	Basic Edition	High-availability Edition
		<ul style="list-style-type: none"> <li>PostgreSQL 9.4</li> <li>PPAS 9.3</li> </ul>
Number of nodes	1	2
Type configuration	Maximum 32-core 128 GB/2 TB	Maximum 60-core 470 GB/3 TB
<i>Set monitoring and alarms</i>	Supported	Supported
<i>IP whitelist</i>	Supported	Supported
Backup and recovery	Supported	Supported
<i>Parameter settings</i>	Supported	Supported
<i>SSL and TDE</i>	Not supported	Supported (Currently MySQL 5.7 does not support TDE)
<i>Log management</i>	Not supported	Supported
Performance optimization	Not supported	Supported
<i>Read-only instance</i> (Additional instances required)	Not supported	Supported only by MySQL 5.6 /5.7
<i>Read/write splitting</i>	Not supported	Supported only in MySQL 5.6
Built-in read/write splitting	Not supported	Not supported
<i>SQL audit</i>	Not supported	Supported (Additional payment required)
High-frequency monitoring	Not supported	Supported (Additional payment required)

## 4.2 Basic Edition

### General introduction

The Basic Edition is a new series of ApsaraDB for RDS, of which the architecture is deployed on a single database node. Compared with the mainstream master/slave High-availability Edition, the Basic Edition contains only one node and has no slave node for fault recovery. Currently, both MySQL and SQL Server support this new series.

The following picture shows the architecture of the Basic Edition and High-availability Edition.



### Comparative advantages

The slave database in the High-availability Edition is used only for failover but does not provide services, and the database duplication adds performance cost to the master database. In this respect, the Basic Edition is not inferior but superior to the High-availability Edition in terms of performance.

The Basic Edition of RDS uses the underlying data distributed storage layer to guarantee the stability of multiple replicas, that is, the fault in or damage to one physical node does not result in data loss. Besides, the costs can be greatly saved by reducing one database node, making the price of the Basic Edition half that of the High-availability Edition.



#### Note:

Because there is only one database node, when the node fails, it takes longer to recover the node. Therefore, the High-availability Edition is recommended for sensitive businesses that have higher requirements on database availability.

### Restrictions

Compared with the High-availability Edition, the Basic Edition does not support the following functions:

- Master/slave switchover
- Zone switch
- Safe connection mode
- Log management
- Performance diagnosis
- Read-only instances

- Disaster tolerance instances

## 5 Instance types

### 5.1 Instance type overview

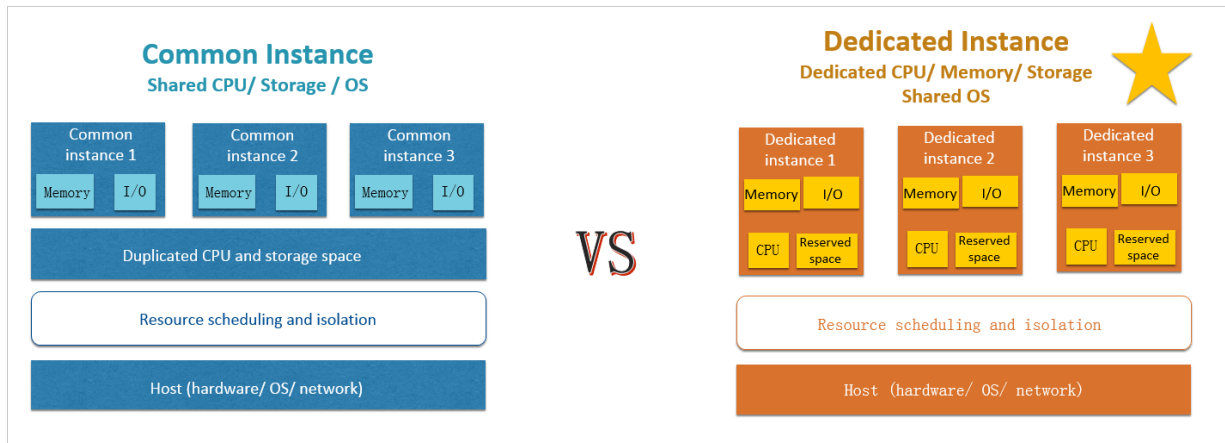
The instances of ApsaraDB for RDS are now available in different types: the common instance, the dedicated instance, and the dedicated host. To distinguish from the instance type available earlier (before January 2017) which we defined as common instance, the new instance type is defined as dedicated instance. The dedicated host can be deemed as the top-level configurations of the dedicated instance. The dedicated host has been so upgraded that it is equivalent to a single full physical machine. Therefore, its resources are dedicated too.

The following list shows the features and applicable scenarios of various instance types:

Type	Description	Applicable scenarios
Common instance	<ul style="list-style-type: none"><li>It is a highly cost-effective instance type which maximizes utilization by reusing resources and allows you to enjoy the benefits of scale.</li><li>Its storage capacity is not linked with CPU/memory , allowing for flexible configurations.</li></ul>	<ul style="list-style-type: none"><li>It is applicable to price-sensitive customers.</li><li>It is applicable for scenarios where high-performance stability is not required.</li></ul>
Dedicated instance/ dedicated host	<ul style="list-style-type: none"><li>It is a new RDS instance type which features fixed computing capabilities , storage space and IO performance.</li><li>A number of fully dedicated CPU cores and threads are assigned to it to ensure long-term stability in computing performance.</li><li>It has a storage space reserved for higher stability .</li><li>The dedicated host is the dedicated instance type</li></ul>	It is applicable for business scenarios where a database -centric system is used, such as finance, e-commerce, government, medium and large -size Internet businesses.

Type	Description	Applicable scenarios
	with top-level configurations	

The differences between the common and dedicated instances are shown in the following figure.



## 5.2 Dedicated instance

### Product introduction

As a new instance type of ApsaraDB for RDS, the dedicated instance features fixed computing capabilities, storage space and IO performance. It differs from other instance types in terms of resource allocation strategies used. With more stable performance, a dedicated instance is the best option for business scenarios where a database-centric system is used, such as finance, e-commerce, government, medium and large -size Internet businesses.

The dedicated instances are available in many types. For more information, see [Instance type list](#). The dedicated host is the top-level configuration of dedicated instance. You can change the configuration of a instance type at will. You are free of restrictions on instance types, which means configuration changes are allowed across different instance types.

### Design principles and performance features

- With the help of the operating system (Linux/Windows) kernel, RDS isolates computing resources on physically different instances. A dedicated instance uses a slightly different CPU allocation strategy from a common instance. A number of fully dedicated CPU cores and threads are assigned to the dedicated instance to ensure long-term stability and predictability in computing performance. This avoids noisy neighborhood on physical machines.
- A dedicated instance has a reserved storage space. As compared with a common instance, the dedicated instance can fully avoid instance migration across physical machines as a result

of the disk capacity increase on your instance or other instances to provide higher stability. Furthermore, the dedicated instance supports hot standby, so you can fail over at any time in case of disk failure on one instance, guaranteeing the availability of the instance. After the fail-over, you can hot swap the faulty host transparently yet unawares.

### Contrast benefits

It is not likely to make direct comparison between dedicated and common instances, because their metrics do not fully match. However, to help make purchase decisions, we select two similar instance types for a cost performance analysis.

Type	Type ID	CPU/ Memory	Disk space	Maximum number of connections	Maximum IOPS	Monthly price (effective before January 2017)
Common instance	rds.mysql.m1.medium	4-core 16 GB	500 GB	4000	7000	2,100 RMB
Dedicated instance	mysql.x8.large.2	4-core 32 GB	500 GB	5000	9000	3,650 RMB

The above table shows that the dedicated instance costs 70% more than the common instance, but it offers twice the memory capacity as the common instance, 25% more in the maximum number of connections, and 28% more in the maximum IOPS. In addition, it has stable CPU computing performance. Therefore, the dedicated instance provides higher overall cost performance in applicable business scenarios. For more information about the prices of the dedicated instances, see [Pricing](#).

## 5.3 Instance type list

### RDS for MySQL

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
Basic series	5.7	Common instance	mysql.n1.micro.1	1 C 1 GB	2000	IOPS =min	20 GB - 1000 GB

			mysql.n2.small. 1	1 C 2 GB	2000	{30* Storage size , 20000}	20 GB - 2000 GB
			MySQL . n2. medium. 1	2 C 4 GB	4000		
			mysql.n4. medium. 1	2 C 8 GB	6000		
			mysql.n4. large. 1	4 C 16 GB	8000		
			mysql.n4. xlarge. 1	8 C 32 GB	10,000		
			mysql.n4. 2xlarge. 1	16 C 64 GB	15,000		
			mysql.n4. 4xlarge. 1	32 C 128 GB	20,000		
			mysql.n8. 4xlarge. 1	32 C 256 GB	64,000		
			mysql.n4. 8xlarge. 1	56 C 224 GB	64,000		
			mysql.n8. 8xlarge. 1	56 C 480 GB	64,000		
High-availability series	5.5/5.6/5.7	Common instance	rds.mysql.t1.small	1 C 1 GB	300	600	5 GB - 2000 GB
			rds.mysql.s1.small	1 C 2 GB	600	1000	
			rds.mysql.s2.large	2 C 4 GB	1200	2000	
			rds.mysql.s2.xlarge	2 C 8 GB	2000	4000	
			rds.mysql.s3.large	4 C 8 GB	2000	5000	
			rds.mysql.m1. medium	4 C 16 GB	4000	7000	
			rds.mysql.c1.large	8 C 16 GB	4000	8000	

			rds.mysql.c1.xlarge	8 C 32 GB	8000	12,000	
			rds.mysql.c2.xlarge	16 C 64 GB	16,000	14,000	5 GB - 3000 GB
			rds.mysql.c2.xlp2	16 C 96 GB	24,000	16,000	
		Dedicated instance (large Memory)	mysql.x8.medium. 2	2 C 16 GB	2500	4500	250 GB
			mysql.x8.large. 2	4 C 32 GB	5000	9000	500 GB
			mysql.x8.xlarge. 2	8 C 64 GB	10,000	18,000	1000 GB
			mysql.x8.2xlarge. 2	16 C 128 GB	20,000	36,000	2000 GB or 3000 GB
		Dedicated instance (more CPUs)	mysql.x4.large. 2	4 C 16 GB	2500	4500	250 GB or 500 GB
			mysql.x4.xlarge. 2	8 C 32 GB	5000	9000	500 GB or 1000 GB
			mysql.x4.2xlarge. 2	16 C 64 GB	10,000	18,000	1000 GB , 2000 GB or 3000 GB
			mysql.x4.4xlarge. 2	32 C 128 GB	20,000	36,000	2000 GB or 3000 GB
		Dedicated Host	rds.mysql.st.d13	30 C 220 GB	64,000	20,000	3000 GB
			rds.mysql.st.h43	60 C 470 GB	100,000	50,000	3000 GB
Finance series (formerly known as Enterprise Multi-node Edition)	5.6	Dedicated Instance (more CPUs)	mysql.x4.large. 3	4 C 16 GB	2500	4500	250 GB or 500 GB
			mysql.x4.xlarge. 3	8 C 32 GB	5000	9000	500 GB or 1000 GB
			mysql.x4.2xlarge. 3	16 C 64 GB	10,000	18,000	1000 GB , 2000 GB

							or 3000 GB
			mysql.x4.4xlarge. 3	32 C 128 GB	20,000	36,000	2000 GB or 3000 GB
		Dedicated instance (large memory)	mysql.x8.medium. 3	2 C 16 GB	2500	4500	250 GB
			mysql.x8.large. 3	4 C 32 GB	5000	9000	500 GB
			mysql.x8.xlarge. 3	8 C 64 GB	10,000	18,000	1000 GB
			mysql.x8.2xlarge. 3	16 C 128 GB	20,000	36,000	2000 GB or 3000 GB
			mysql.x8.4xlarge. 3	32 C 256 GB	40,000	72,000	3000 GB
		Dedicated Host	mysql.st.8xlarge. 3	60 C 470 GB	100,000	120,000	3000 GB
Read-only instance	5.6/5.7	Common instance	rds.mysql.t1.small	1 C 1 GB	300	600	5 GB - 2000 GB
			rds.mysql.s1.small	1 C 2 GB	600	1000	
			rds.mysql.s2.large	2 C 4 GB	1200	2000	
			rds.mysql.s2.xlarge	2 C 8 GB	2000	4000	
			RDS. MySQL.s3.large	4 C 8 GB	2000	5000	
			rds.mysql.m1.medium	4 C 16 GB	4000	7000	
			rds.mysql.c1.large	8 C 16 GB	4000	8000	
			rds.mysql.c1.xlarge	8 C 32 GB	8000	12,000	

			rds.mysql.c2.xlarge	16 C 64 GB	16,000	14,000	5 GB - 3000 GB
			rds.mysql.c2.xlp2	16 C 96 GB	24,000	16,000	
		Dedicated instance (large memory)	mysqlro.x8.medium.1	2 C 16 GB	2500	4500	250 GB
			mysqlro.x8.large.1	4 C 32 GB	5000	9000	500 GB
			mysqlro.x8.xlarge.1	8 C 64 GB	10,000	18,000	1000 GB
			mysqlro.x8.2xlarge.1	16 C 128 GB	20,000	36,000	2000 GB or 3000 GB
		Dedicated instance (more CPUs)	mysqlro.x4.large.1	4 C 16 GB	2500	4500	250 GB or 500 GB
			mysqlro.x4.xlarge.1	8 C 32 GB	5000	9000	500 GB or 1000 GB
			mysqlro.x4.2xlarge.1	16 C 64 GB	10,000	18,000	1000 GB , 2000 GB or 3000 GB
			mysqlro.x4.4xlarge.1	32 C 128 GB	20,000	36,000	2000 GB or 3000 GB
		Dedicated Host	rds.mysql.st.d13	30 C 220 GB	64,000	20,000	3000 GB

**RDS for SQL Server**

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
Basic series	2012 Enterprise (formerly known as 2012)	Common instance	rds.mssql.s2.large	2 C 4 GB	Not limited	IOPS =min {30* Storage size , 20000}	20 GB - 2000 GB
			rds.mssql.s2.xlarge	2 C 8 GB			
			rds.mssql.s3.large	4 C 8 GB			
			rds.mssql.m1.medium	4 C 16 GB			
			rds.mssql.c1.large	8 C 16 GB			
			rds.mssql.c1.xlarge	8 C 32 GB			
			rds.mssql.c2.xlarge	16 C 64 GB			
	2012/2016 Web	Dedicated instance	mssql.x2.medium.w1	2 C 4 GB	Not limited	IOPS =min {30* Storage size	20 GB - 2000 GB

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
			mssql.x2.large.w1	4 C 8 GB		20000}	
			mssql.x2.xlarge.w1	8 C 16 GB			
			mssql.x2.2xlarge.w1	16 C 32 GB			
			mssql.x4.medium.w1	2 C 8 GB			
			mssql.x4.large.w1	4 C 16 GB			
			mssql.x4.xlarge.w1	8 C 32 GB			
			mssql.x4.2xlarge.w1	16 C 64 GB			
High-availability series	2008 R2 Enterprise	Common instance	rds.mssql.s1.small	1 C 2 GB	600	1000	10 GB - 2000 GB
			rds.mssql.s2.large	2 C 4 GB	1200	2000	
			rds.mssql.s2.xlarge	2 C 8 GB	2000	4000	
			rds.mssql.s3.large	4 C 8 GB	2000	5000	
			rds.mssql.m1.medium	4 C 16 GB	4000	7000	
			rds.mssql.c1.large	8 C 16 GB	4000	8000	

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
			rds.mssql. c1.xlarge	8 C 32 GB	8000	12,000	
			rds.mssql. c2.xlarge	16 C 64 GB	16,000	14,000	
			rds.mssql. c2.xlp2	16 C 96 GB	24,000	16,000	
		Dedicated instance	mssql.x8. medium. 2	2 C 16 GB	2500	4500	250 GB
			mssql.x8. large. 2	4 C 32 GB	5000	9000	500 GB
			mssql.x8. xlarge. 2	8 C 64 GB	10,000	18,000	1000 GB
			mssql.x8. 2xlarge. 2.	16 C 128 GB	20,000	36,000	2000 GB
		Dedicated host	rds.mssql. st.d13	30 C 220 GB	64,000	20,000	2000 GB
			rds.mssql. st.h43	60 C 470 GB	100,000	50,000	2000 GB

### RDS for PostgreSQL

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
High- availability series	9.4	Common instance	rds.pg.t1. small	1 C 1 GB	100	600	5 GB - 2000 GB
			rds.pg.s1. small	1 C 2 GB	200	1000	

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
			rds.pg.s2.large	2 C 4 GB	400	2000	
			rds.pg.s3.large	4 C 8 GB	800	5000	
			rds.pg.c1.large	8 C 16 GB	1500	8000	
			rds.pg.c1.xlarge	8 C 32 GB	2000	12,000	
			rds.pg.c2.xlarge	16 C 64 GB	2000	14,000	
		Dedicated instance (large memory)	pg.x8.medium. 2	2 C 16 GB	2500	4500	250 GB
			pg.x8.large. 2	4 C 32 GB	5000	9000	500 GB
			pg.x8.xlarge. 2	8 C 64 GB	10,000	18,000	1000 GB
			pg.x8.2xlarge. 2	16 C 128 GB	12,000	36,000	2000 GB
		Dedicated instance (more CPUs)	pg.x4.large. 2	4 C 16 GB	2500	4500	250 GB or 500 GB
			pg.x4.xlarge. 2	8 C 32 GB	5000	9000	500 GB or 1000 GB
			pg.x4.2xlarge. 2	16 C 64 GB	10,000	18,000	1000 GB or 2000 GB
			pg.x4.4xlarge. 2	32 C 128 GB	12,000	36,000	2000 GB or 3000 GB
		Dedicated Host	RDS. PG. St. D13	30 C 220 GB	4000	20,000	3000 GB

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
			rds.pg.st. h43	60 C 470 GB	4000	50,000	3000 GB

**RDS for PPAS**

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
High-availability series	9.3	Common instance	rds.ppas.t1.small	1 C 1 GB	100	600	5 GB - 2000 GB
			rds.ppas.s1.small	1 C 2 GB	200	1000	
			RDS.PPAs.s2.large	2 C 4 GB	400	2000	
			rds.ppas.s3.large	4 C 8 GB	800	5000	
			rds.ppas.m1.medium	4 C 16 GB	1500	8000	
			rds.ppas.c1.xlarge	8 C 32 GB	2000	12,000	
			rds.ppas.c2.xlarge	16 C 64 GB	2000	14,000	
		Dedicated instance	ppas.x8.medium. 2	2 C 16 GB	2500	4500	250 GB
			ppas.x8.large. 2	4 C 32 GB	5000	9000	500 GB

Series	Version	Type	Type code	CPU/ Memory	Maximum number of connections	Maximum IOPS	Storage size
			ppas.x8.xlarge. 2	8 C 64 GB	10,000	18,000	1000 GB
			ppas.x8.2xlarge. 2	16 C 128 GB	12,000	36,000	2000 GB
		Dedicated Host	rds.ppas.st.d13	30 C 220 GB	4000	20,000	3000 GB
			rds.ppas.st.h43	60 C 470 GB	4000	50,000	3000 GB

### History instance type

The history instance type of RDS for SQL Server is in the following table. It's not provide for creating the new RDS instance. Recommend the user to use the latest specifications.

Type code	CPU(Core)	Memory (MB)	Maximum number of connections	Maximum IOPS
rds.mys2.small	2	240 MB	60	150
RDS. mys2.mid	4	600 MB	150	300
rds.mys2.standard	6	1200 MB	300	600
rds.mys2.large	8	2400 MB	600	1200
rds.mys2.xlarge	9	6000 MB	1500	3000
rds.mys2.2xlarge	10	12,000 MB	2000	6000
rds.mys2.4xlarge	11	24,000 MB	2000	12,000
rds.mys2.8xlarge	13	48,000 MB	2000	14,000

### RDS for SQL Server

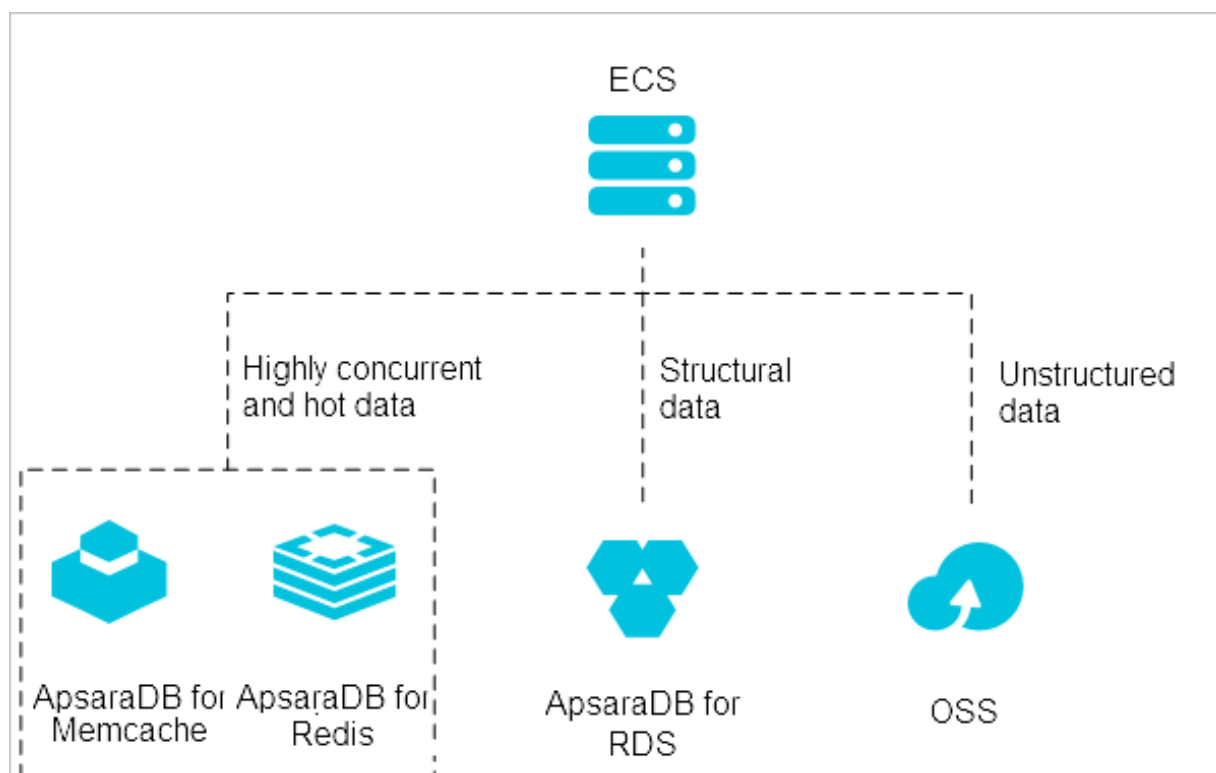
The history instance type of RDS for SQL Server is in the following table. It's not provide for creating the new RDS instance. Recommend the user to use the latest specifications.

Type code	CPU(Core)	Memory (MB)	Maximum number of connections	Maximum IOPS
rds.mss1.small	6	1000 MB	100	500
rds.mss1.mid	8	2000 MB	200	1000
rds.mss1.standard	9	4000 MB	400	2000
rds.mss1.large	10	6000 MB	600	3000
rds.mss1.xlarge	11	8000 MB	800	4000
rds.mss1.2xlarge	12	12,000 MB	1200	6000
rds.mss1.4xlarge	13	24,000 MB	2000	12,000
rds.mss1.8xlarge	13	48,000 MB	2000	14,000

## 6 Typical applications

### 6.1 Diversified data storage

RDS supports diversified storage extension through Memcache, Redis, and OSS storage products



#### Cache data persistence

ApsaraDB for Memcache and ApsaraDB for Redis can be used together to form a high-throughput and low-latency storage solution,

with a request delay of

- only a few milliseconds
- and a cache that supports a higher QPS (queries per second) than standard RDS.

For more information, see [#####](#).

#### Multi-structure data storage

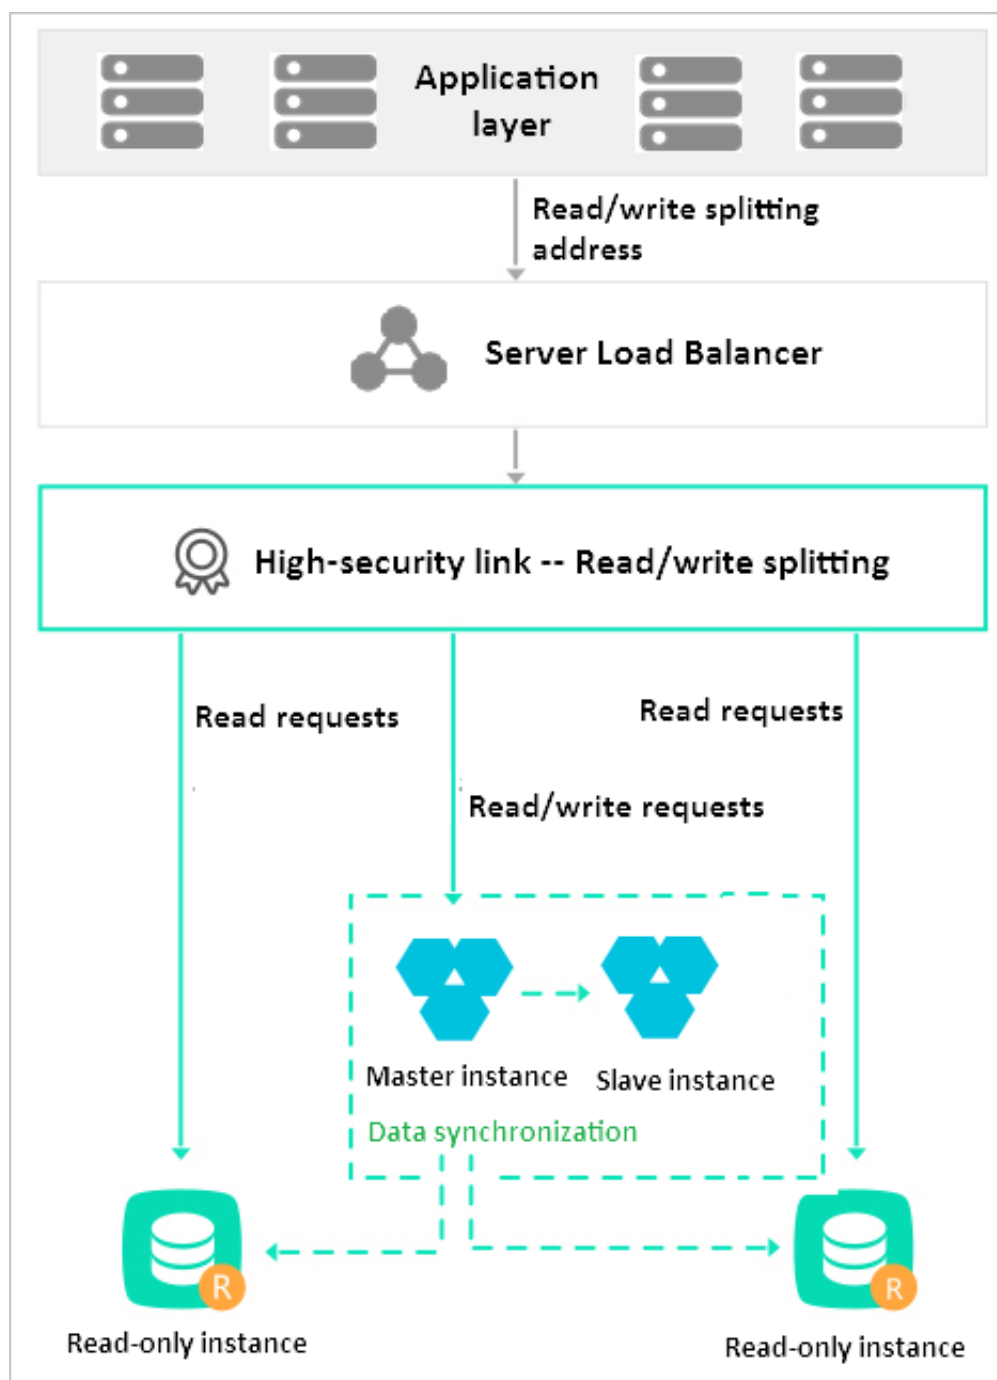
OSS is an Alibaba Cloud storage service that features massive capacity, robust security, low cost, and high reliability. RDS and OSS can work together to form multiple data storage solutions.

For example, when RDS and OSS are used in a forum, resources such as the images of registered users or those posted on the forum can be stored in OSS to reduce the storage pressure on RDS.

For more information, see [#####](#).

## 6.2 Read/write splitting function

ApsaraDB for MySQL allows read-only instances to be directly attached to RDS in order to distribute the read pressure on the master instance. MySQL Each read-only instance has an independent connection string, and the read pressure can be automatically distributed on the application side.



For more information on creating read-only instances on ApsaraDB for MySQL, refer to #####.

## 7 RDS usage instructions

---

- Instructions for RDS instance upgrading

There is a transient disconnection of up to 30 seconds during RDS instance upgrading. To avoid glitches and served downtime, you must prepare for the flash in advance and set the program for auto-reconnection to RDS.

- Risks of intranet/Internet switch

In the process of switching between intranet and Internet, RDS may be disconnected with the server and the RDS IP address may change. After the switch, update the connection address in the program in a timely manner.

- Rollback risks

Currently, RDS only supports data rollback for the entire instance, not for a single table or database. It is important to make backups of crucial data before rollback to prevent data loss. If you need to roll back only partial tables or partial data, we recommend that you recover the data by creating a temporary instance or a clone instance, once a temporary instance or clone instance is created successfully, the required data is directed back to the production library, see for details.

- Instructions on RDS locking policies

Currently, RDS only supports data rollback for the entire instance, not for a single table or database. It is important to make backups of crucial data before rollback to prevent data loss. We recommend that you create a temporary instance for data recovery if you only want to roll back part of the tables or part of the data. You must import the required data back to the production databases from the temporary instance. For more information, see [#####](#).

- Instructions on RDS failover

RDS adopts a master-slave mode to guarantee high availability. When the master instance fails, RDS switches to the slave instance within 30 seconds. In the failover process, there may be an inaccessible period of (less than or equal to) 30 seconds. You must set auto-reconnection of your program with RDS to avoid the service downtime.

- Instructions on data synchronization mode for MySQL instances

— For MySQL 5.1, data synchronization between the master and slave instances adopts an asynchronous mode. This mode guarantees high performance of instances, but it may lead to probabilistic data inconsistency between the master and slave instances.

- For MySQL 5.5, data synchronization between the master and slave instances adopts a semi-synchronous mode. This mode compromises write performance to some extent, but it can greatly reduce the probability of data inconsistency between the master and slave databases. If you have a high requirement on data reliability (for example, financial systems), we recommend that you purchase MySQL 5.5 instances or later versions.
- For MySQL 5.6, data synchronization between the master and slave instances adopts GTID (a new feature in MySQL 5.6). This feature guarantees both the performance and data consistency.
- Notes of attentions for using RDS

After purchasing an RDS instance, you do not need to conduct basic O&M for the database (such as high availability, backup, and security patches). But you must pay attention to the following matters:

- Check whether the CPU, IOPS, space, and the number of connections of your RDS instance are adequate. If not, you must optimize or upgrade them.
- Check whether your RDS instances have performance problems, whether there is a large number of slow SQL queries, whether the SQL statements need optimization, and whether there are excess or missing indexes.
- Check whether your RDS instances have any web SQL injection warnings. If so, your database may be vulnerable to web SQL injections. You must modify the application to prevent web SQL injection attacks.