Alibaba Cloud
Apsara Stack Agility SE

Glossary

Version: 1912, Internal: V3.1.0

Issue: 20200312

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.
5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified,

reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

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

Document conventions

Style	Description	Example
	A danger notice indicates a situation that will cause major system changes , faults, physical injuries, and other adverse results.	 **Danger:** Resetting will result in the loss of user configuration data.
	A warning notice indicates a situation that may cause major system changes , faults, physical injuries, and other adverse results.	 **Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
	A caution notice indicates warning information, supplementary instructions , and other content that the user must understand.	 **Note:** If the weight is set to 0, the server no longer receives new requests.
	A note indicates supplemental instructions, best practices, tips, and other content.	 **Note:** You can use **Ctrl** + **A** to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings** > **Network** > **Set network type**.
Bold	Bold formatting is used for buttons, menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands.	Run the `cd /d C:/window` command to enter the Windows system folder.
`Italic`	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *`Instance_ID`*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig` *`[-all|-t]`*
{} or {a|b}	This format is used for a required value , where only one item can be selected.	`switch` *`{active|stand}`*

Contents

Glossary

A | B | C | D | E | F | G | H | I | K | L | M | N | O | P | R | S | T | U | V | W

A

AnalyticDB for PostgreSQL

AnalyticDB for PostgreSQL is a data warehousing service that provides online Massively Parallel Processing (MPP). AnalyticDB for PostgreSQL is developed based on the Greenplum Open Source Database program and is enhanced with some in-depth extensions by Alibaba Cloud. It supports features including the OSS external table, JSON data type, and HyperLogLog estimation analysis. AnalyticDB for PostgreSQL provides flexible hybrid analysis functions through query and anticipation of SQL2008 standard and OLAP analytic and aggregate function; it also supports the hybrid storage of row storage and column storage, data compression technology to reduce the storage cost. AnalyticDB for PostgreSQL provides online scalability, backup, performance monitoring, and other services. Users do not have to manage complicated large-scale MPP clusters. AnalyticDB for PostgreSQL allows DBA, developers, and data analysts to focus on how to enhance productivity through SQL and create enterprise core values.

Application Migration to the Cloud

Application Migration to the Cloud: specific to the user requirements, reconstructs the in-cloud design and the applications in between data layers and middleware layers, to complete the system migration to and operation in the cloud.

ApsaraDB for PostgreSQL

PostgreSQL is one of the most advanced open-source database. As the ancestor of the academic relational database management system, ApsaraDB for PostgreSQL features by the complete implementation of the SQL rules and abundant support of data types, including JSON data, IP data, geometric data, and etc. Most of these supported data types are not supported by most other commercial database. Except for functions like complete support of tasks, sub queries, multi-version control, data integrity check, ApsaraDB for PostgreSQL is integrated with key features like high availability and backup/restore to reduce the operation pressure. Current RDS for PostgreSQL also supports version 9.4.

ApsaraDB for PPAS

A Postgres Plus Advanced Server (PPAS)-based database service. PPAS is a stable, secure, and scalable enterprise-class relational database. It is also the world's most advanced open-source database based on PostgreSQL. It delivers enhanced performance, application solutions, and compatibility, and provides the capability to run Oracle applications directly. You can run enterprise-class applications on PPAS stably and obtain cost-effective services. ApsaraDB for PPAS incorporates a number of advanced functions including account management, resource monitoring, backup recovery, and security control, and is updated and improved regularly. RDS for PPAS currently supports version 9.3.

ApsaraDB for RDS

An on-demand, stable, reliable, and scalable online relational database service. It has multiple security measures and perfect performance monitoring systems and provides professional database backup, recovery and optimization solutions. It enables you to focus on application development and business growth.

ApsaraDB for SQL Server

One of the earliest commercial database systems. It offers excellent performance for complex SQL queries. It fully supports .NET-based applications in Windows.

Aurora

The HA (high-availability) management system of Relational Database Service (RDS). It checks the status of the master database every three seconds. When it detects that the master database is down, it switches the SQL requests of users to the slave database quickly.

Aurora

Average Response Time

The average response time, in milliseconds, when DRDS executes logical SQL queries during a cycle. It is calculated as follows: (the time point recorded by the DRDS for the last package of the result set) - (the time point at which DRDS receives the SQL query).

B

Back-end server

A set of cloud servers that accept load-balanced distribution requests. Server Load Balancer forwards access requests to this set of servers based on the rules you set.

Binary Log File Position

A method used in MySQL binary log files. Each line records a data change action. The position of the line is called the binary log (binlog) file position.

Bucket

A unit that is used to manage stored objects. Objects must belong to a specific bucket. A container that is used to store objects. ZStack uses the buckets in Object Storage Service (OSS) to upload image files.

Buffer Pool

A cache of existing database connections. It allows applications to reuse database connections that already exist in the pool, and reduces repeated attempts to create new database connections. This technology can improve system performance and avoid the overhead of building a new connection.

C

Classless Inter-Domain Routing

A method used to assign IP addresses to users and efficiently route IP packets on the Internet to classify the IP addresses.

CIDR

Connection Mode

A mode that defines the transmission path of data packages for RDS instances. Performance is the standard connection mode. Safety is the high security connection mode

Cookie

Data (usually encrypted) that websites store on your local terminal for identity identification and session tracking.

D

Data Compression

A technique that uses algorithms to reorganize data, reduce redundant data, and save storage space without loss of valid data. It can improve the data transmission, storage, and processing efficiency.

Data Consumed Time Stamp

A time stamp that records the last time subscription data was consumed by the downstream SDK. The SDK sends an ACK packet to the DTS server for each piece of consumed data. The DTS server updates and saves the respective time stamp. When the SDK restarts unexpectedly, the DTS server automatically pushes subscription data based on the latest time stamp.

Data Migration to the Cloud

Migration of data from traditional architectures to the cloud.

Data Scope

The scope of the time stamp for the incremental data stored in the prescription channel. The time stamp of the incremental data corresponds to the time stamp when the piece of incremental data is applied and recorded by event log in the RDS instance. By default the prescription channel only stores the incremental data of one day. The data transmission service cleans the incremental data regularly, and update the data scope of prescription channel.

Data Transmission Service

A service that helps migrate data between different data sources such as relational databases, NoSQL, and OLAP. DTS integrates data migration, data subscription, and real-time data synchronization. It is provides millisecond-level long-distance asynchronous data transmission in public cloud and hybrid cloud scenarios. Its underlying data stream uses the remote active-active infrastructure used for Double 11. The infrastructure has provided stable real-time data streaming for thousands of downstream applications for three years. You can use DTS to easily create secure, scalable, and highly available data architecture.

DTS

Data Update

A type of database update. Data updates only modify data. They do not modify the schema.

Database Engine

The core service used to store, process, and protect data. It provides controlled access and rapid transaction processing to meet the requirements of enterprise applications that process massive data. It can be used to create relational databases that process transactions and data analysis online. This includes the creation of tables that store data and database objects that are used to view, manage, and guarantee data security (such as indexes, views, and storage procedures).

Destination Network Address Translation

The action of translating the destination address of an IP packet to another address.

DNAT

Distributed Data Access Engine

A method that is used to manage distributed databases. Taobao Distributed Data Layer (TDDL) is a distributed data access engine that is built on the unified configuration of JDBC Data source. TDDL is used for data replication between heterogeneous databases. TDDL ensures that users are not aware of database and table splitting. TDDL supports features such as the active-standby mode, read/write splitting, and dynamic database configuration.

Distributed Relational Database Service

Distributed Relational Database Service (DRDS) is a lightweight, flexible, stable, and efficient middleware product developed by Alibaba Cloud. DRDS focuses on expanding standalone relational databases, and has been tested by core transaction links in Tmall, such as during the Singles' Day Shopping Festival. DRDS has been used for ten years and is a trusted database service provider.

DRDS

Domain Name

The name of a computer or a computer group on the Internet. It consists of several parts separated by periods (.) to identify the electronic location of the computer or computer group for data transfer.

DRDS Hint

A custom hint provided by DRDS to specify certain actions. DRDS Hint uses relevant syntax to affect and optimize the way SQL statements are executed.

DRDS Instance

An instance that consists of several DRDS servers. Multiple DRDS databases can be stored in a DRDS instance.

DRDS Instance ID

A unique ID that identifies a DRDS instance.

DRDS Note

Custom comments provided by DRDS to specify special behavior such as specific syntax that affects the way SQL is executed.

DRDS Sequence

A 64-bit number that corresponds to the BIGINT type in MySQL. It is used to generate a globally unique number sequence typically used to generate primary key columns and unique index columns.

DRDS Server Count

The number of DRDS servers in a DRDS instance.

E

Endpoint

The access domain name served by the OSS over the Internet or intranet.

Exclusive Virtual Host

Exclusive Virtual Host is the complete server resource with no fight for the resource, which makes it more stable,with no limit of traffic, which make it more speedy, with independent IP, which makes it easier to promote. It suits more for the enterprise website project.

Existing Data Synchronization

Initialization of the historical data of the synchronization object to the target instance before the link incremental data is synchronized.

F

Forwarding Weight

The weight for traffic forwarding that you can specify for each Elastic Compute Service (ECS) instance in the backend server pool according to the external service capabilities and situations of the backend ECS instance. The higher the weight is, the more access requests are assigned to the ECS instance.

Full Backup

The process of backing up all data or applications at a time point.

Full Data Migration

A type of migration task. It refers to the migration of all data except Data Definition Language (DDL) from the source instance database to the target instance. If you select only full data migration but not incremental data migration when you create a migration task, the data newly added to the source instance will not be migrated to the target instance during the migration process.

Full Migration

A mode of data migration. Within the specified period, all data is migrated from the source system to the target system in a way data consistency is guaranteed.

Full Table Scan

An action in database sharding mode where DRDS executes an SQL statement on all sub-tables and then merges the results if the SQL statement does not contain the shard key. We recommend that you do not perform this action because it negatively affects database performance.

Function Compute

An event-driven serverless computing service. Users only need to write and upload the code without the need to manage infrastructure such as hosts. Function Compute prepares the computing resources for you and runs your code in an elastic and reliable manner.

G

Globally Unique ID

A globally unique numerical ID. It helps keep data such as unique keys and primary keys globally unique in a distributed environment.

GUID

H

Health Check

A basic function of Server Load Balancer, which is used to check the business availability of the back-end servers (Elastic Compute Service (ECS) instances), improve the overall availability of the front-end business, and avoid the effect on the overall services caused by the abnormal back-end ECS instances.

Heterogeneous Indexing

A solution provided by DRDS to perform a full table scan if no shard key is available. For partition tables in DRDS, if shard keys are available in WHERE conditions, WHERE conditions in SQL statements are routed to a specific sub-database to improve query efficiency.

High-Availability Virtual IP Address (HAVIP)

An intranet IP resource that can be created and released independently. You can use the ARP protocol on the Elastic Compute Service (ECS) instances to declare the IP.

HAVIP

Horizontal Sharding

The partitioning of table columns into multiple tables based on the specified rule to achieve horizontal extension.

HyperText Transfer Protocol

A universal, stateless, and application-oriented protocol for distributed, collaborative, and hypermedia information systems.

HTTP

I

Image Processing Service (IMG)

An extended service of Object Storage Service (OSS). It is built on OSS and provides real-time image processing functions, such as resizing, cropping, reformatting, and adding watermarks to images. You must enable image processing in the OSS console first, and then call the RESTful interface provided by the image processing service to process the images stored on OSS, and obtain the processed images immediately.

Imported Database

A database that is created in the MySQL instance and mounted to DRDS.

Incremental Backup

A type of backup that only copies files that have changed since the previous backup.

Incremental Data Migration

A type of data migration. It synchronizes the incremental data from the source instance to the target instance during the migration process. If you select both full data migration and incremental

data migration when you create a migration task, DTS first performs static snapshot on the source instance, migrates the snapshot data to the target instance, and then synchronizes the incremental data written to the source instance to the target instance during the migration process. Incremental data migration is a process of real-time data synchronization between the source and target instances. It does not end automatically. If you want to terminate the migration, terminate it manually on the console.

Incremental Migration

A type of data migration. It synchronizes part of data changes (addition, deletion, and modification) in a certain period from the source system to the target system to ensure data consistency. In some systems, transactional tables are locked during dynamic incremental migration, causing a data write failure during the migration. Exercise caution when using dynamic incremental migration.

InnoDB Buffer Pool

InnoDB keeps a buffer pool in its memory for data buffer and index. The buffer pool is divided into two sections, section of sublist of new blocks (frequently visited data-hot data), section of sublist of old blocks (infrequently visited data). When the user visits data and the data is available from buffer pool, it is returned directly, otherwise the disk data will be loaded to the buffer pool of sublist of old blocks section first and then move to the sublist of new blocks seciton. The data from sublist of old blocks section get flushed according to its access frequency through LRU.

IO Volume

The amount of data read and written from the database per second.

K

Keepalived

An open-source software that guarantees the high availability of clusters and detects the status of Web servers. If it detects that a Web server is down or has an error, the Web server is automatically removed from the system. The removed Web server is automatically added after resuming the normal work.

L

Layer-4 Server Load Balancer

A service that uses the information defined at the network transport layer (layer 4) as the basis for deciding how to distribute client requests across a group of servers. For Internet traffic, it makes the load balancing policy based on the source and destination IP addresses and ports recorded in the packet header, without considering the contents of the packet.

Layer-7 Server Load Balancer

A service that provides load balancing for services at the application level, such as HTTP services. It determines the routing based on the characteristics of the HTTP header and the actual contents of the message, such as the URL, data type (text, video, and graphics), and information in the cookie.

Linux Virtual Server

A cluster technology that supports the IP load balancing and content-based request distribution techniques.

LVS

Listener

A concept used in Server Load Balancer instances. A listener defines how the incoming requests are distributed. You must add at least one listener to a Server Load Balancer instance. It includes frontend ports, backend ports, Server Load Balancer policies, and health check configurations. Each listener corresponds to a backend application service.

Live Migration

Also known as real-time migration. It refers to the process of migrating a running VM or application between different hosts without disrupting its normal operation. Its logic migration procedure is almost the same as that of offline migration. The difference is that live migration results in a very short downtime to ensure that VMs can provide services normally. During the early stage of migration, the services run on the source host. When the migration reaches a certain stage, the target host has the necessary resources to run the system. After a transient switchover, the source host takes services over from the target host, ensuring service continuity. For services, users are not aware of the service interruption because the switchover is transient. Live migration applies to scenarios that require high service availability.

Locking Mechanism

The locking mechanism of ApsaraDB for RDS. If the storage space of your instance is full, the system will lock your instance into the read-only status.

Logical Backup

The process of using SQL statements to extract data from databases and then saving it in binary files.

Logical SQL

An SQL query sent from an application to DRDS.

M

Mass Storage

The storage of large amounts of data in a persisting fashion.

Massively Parallel Processing

A distributed shared-nothing computing architecture, where each node is completely independent of other nodes and multiple nodes perform concurrent computing to improve performance.

MPP

Multipart Upload

A method of uploading a large file by separating it into smaller parts.

Multi-zone

An area composed of multiple zones, which is deployed in different regions by Alibaba Cloud Server Load Balancer to achieve disaster recovery across data centers in the same region. If the data center in the active zone malfunctions or becomes unavailable, Server Load Balancer rapidly switches to the data center in the standby zone to restore its service capabilities within 30s. When the active zone is restored, Server Load Balancer automatically switches back to the data center in the active zone to provide services.

N

Non-sharding Mode

Refers to the use of a database has been built in MySQL(RDS) as a DRDS database. In this mode the library of DRDS sub-table and other functions are not available, and only the DRDS read and write separate functions are available.

Number of Active Threads

The number of threads that DRDS is running. It can be typically used to represent the current load of DRDS.

O

Object

A basic unit for data storage in OSS. An object is a file in OSS. An object consists of the meats data (Object Meta), user data (data), and object name (key). Objects are identified by a key. The key is unique within the bucket.

Object Storage Service

A secure, reliable, and cost-friendly cloud storage service provided by Apsara Stack to process a large amount of data. You can call APIs to upload or download data to or from any application anytime anywhere, and perform simple data management through the Web console. OSS can store files of any type from various websites, development-oriented enterprises, and developers.

OSS

P

Persistent Connection

A connection used to send multiple data packets continuously. During the connection, both parties must send link detection packets if no data packets are sent.

Physical Average Response Time

The average response time, in milliseconds, when DRDS performs logical SQL queries during a cycle. The response time for an SQL query is calculated based on the formula: Response time for an SQL query = (The point in time recorded by DRDS for the last package of the result set) – (The point in time at which DRDS receives the SQL query).

Physical QPS

The number of physical SQL queries per second that DRDS executes on ApsaraDB for RDS.

Physical SQL

The SQL query that is sent to the RDS after the DRDS parses a logical SQL query.

Pre-check

A mandatory phase before the start of a migration task. It checks the preconditions that will affect the migration, such as connectivity of the source instance, the permissions of the migrating account, etc. If it fails, make the appropriate changes and try again.

Private-Read-Write

One of the three access permissions to Object Storage Service (OSS) buckets and objects.
The other two access permissions are public-read-write and public-read. The private-read-write permission only allows the bucket creator to read and write files in the bucket, while other users cannot access the files in this bucket.

Public-Read

An access control permission level for OSS buckets and objects. The access control permissions include public-read-write, public-read, and private. Public-read allows only the creator of the bucket to write objects to the bucket. Anyone else (including anonymous users) can only read objects from the bucket.

Public-Read-Write

An access control permission level for OSS buckets and objects. The access control permissions include public-read-write, public-read, and private. Anyone (including anonymous users) can read and write objects from the bucket. Any costs incurred by these operations are borne by the creator of the bucket. Exercise caution when using this permission.

R

Read Policy

A policy that defines the ratio of read-only SQL statements sent by DRDS between the primary instance and read-only instance.

Read Weight Proportioning

The proportional load balancing of read-only instances in RDS to achieve real-time read throttling.

Read/Write Splitting

A way to improve DRDS performance. If physical database instances include read-only instances, you can configure the ratio of the read-only SQL statements between the primary and read-only instances on the DRDS console. DRDS automatically identifies the type of SQL statements and runs the SQL statements based on the ratio.

Read-Only Instance

An instance type in which data in the instance can be accessed but not modified. There are two types of physical database types based on whether the write is allowed: the primary instance that allows read and write and the read-only instance that only allows read. The primary instance applies to the primary MySQL instance in Apsara Stack and primary RDS instance in Alibaba Cloud public cloud. The read-only instance applies to the secondary MySQL instance in Apsara Stack and RDS instances in Alibaba Cloud public cloud.

Read-Only SQL

A type of SQL statement for reading data, such as SELECT statements.

Received Traffic

SQL network traffic sent from applications to DRDS.

Recovery Point Objective

The amount of data at risk, which is determined by the amount of time between data protection events and reflects the amount of data that potentially could be lost during a disaster recovery.

RPO

Relational Database Service

A stable, reliable, and elastically scalable online database service. Based on Apsara distributed system and high-performance storage, Relational Database Service (RDS) supports the MySQL engine. It provides a complete set of solutions for disaster recovery, backup, recovery, monitoring, and migration to free you from worries about database operations.

RDS

Round-robin Scheduling

A scheduling algorithm that Server Load Balancer uses to distribute the traffic. Requests are sequentially distributed to the back-end servers.

RR

Route Computing

A computing method that reduces the routing overhead to improve routing protocol functions.

Route Entry

Each item in the route table, which defines the next hop address for the network traffic destined to a specific destination Classless Inter-Domain Routing (CIDR) block.

Route Table

A list of route entries in a VRouter.

RPS

The number of SQL statements performed each second, including begin, commit, insert, delete, and create table.

S

Schema Migration

A type of migration. In database migration, it refers to the migration of schema definitions including tables, views, triggers, stored procedures, and storage functions. For migration between heterogeneous databases, the data type mapping is carried out in the schema migration phase, and the object type is changed to the appropriate target type. For example, when Oracle is migrated to MySQL, the number type in Oracle is mapped to the decimal type in MySQL.

Schema Update

An update that modifies the schema definition such as create table, alter table, drop view and so on. Users can choose whether to subscribe to schema updates when creating a subscription channel.

Sequential Partitioned Message

A message that is partitioned by the sharding key for a topic. Messages in one partition are published and consumed in the strict First In, First Out (FIFO) order. A sharding key is a key field for distinguishing segments of ordered messages. The sharding key is different from the key of a standard message.

Server Load Balancer

A concept of Server Load Balancer instances, which includes front-end ports, back-end ports, Server Load Balancer policies, and health check configurations. Each listener corresponds to a back-end application service.

Server Load Balancer Monitoring

A function of listening to the ports, load balancing policies, and health check configurations.

Session Persistence

A basic function of Server Load Balancer that distributes access requests from the same client to the same backend server for processing.

Shard Key

A column in logical tables, according to which DRDS routes data and SQL statements to physical tables.

Sharding/Partitioning Mode

A mode that allows full use of all DRDS features. It creates multiple libraries for a MySQL instance, and creates a DRDS database from these tables.

Shift

An operation that writes the remaining incremental records from MySQL binlog into DRDS. It is a part of the data import and seamless scaling process.

Short-lived Connections

A connection created when both parties have data exchanges. The connection is closed after the data is sent. The data transmission occurs only once for each connection.

SLB Instance

An instance running the Server Load Balancer (SLB) service.

Small Table Copy

A solution that enables users to implement peer-to-peer synchronization of tables that do not change frequently, such as the configuration and constant tables across nodes. This speeds up the join queries between those tables and other split tables.

Smooth Scaling

DRDS supports adding storage instance nodes to resize the database. Resizing without affecting the access to the original data is called smooth scaling.

Software Defined Network

A method that separates the control level of network devices from the data level to flexibly control the network traffic and provide a good platform for the innovation of core networks and applications.

SDN

Specification

The measure of the performance of different link specifications to the number of simultaneous records.

Static Website

A website with web pages composed of fixed content, including scripts (such as JavaScript) that are executed by the client without the need for prior web programming or database design.

Static Website Hosting

Users can set their own bucket to the static website hosting mode. After the static website hosting mode is configured, the bucket can be accessed as a static website, and users can be automatically redirected to the index page or error page.

Static/dynamic Resource Separation

Separation of static and dynamic resources from Web pages or mobile applications.

Storage

A service that is used to store information or data in a specific format in a specific medium built in or out of a computer. Meanwhile, valid access to the medium is guaranteed.

Storage Engine

A component of the MySQL database that manages how data is stored in files or memory. Each storage technology adopts a different storage mechanism, indexing technique, and locking mechanism. A storage technology provides a unique set of features and capabilities. You can select different storage technologies to obtain additional features such as improved data processing capabilities, improving the overall performance of your applications.

Subscription Channel ID

A unique identification of a subscription channel. After you purchase a subscription channel, Data Transmission Service (DTS) automatically generates a subscription channel ID. When you use the SDK to consume incremental data, you need to configure the appropriate subscription channel ID. The subscription channel ID corresponding to each subscription channel is displayed in the subscription list of the DTS console.

Sub-database

The physical database of each storage instance. Data in the logical database is stored in multiple physical storage instances after the DRDS database is split horizontally.

Synchronization Latency

The difference between the time stamp of the target instance in the source library, and the current time stamp of the homologous instance.

Synchronized Records per Second

The number of records synchronized to the target instance per second (unit: RPS). Synchronized RPS is the specification by which the data synchronization service is sold.

T

Table Broadcast

An action in which data of a single-database table is synchronized in advance to all sub-libraries. DRDS then converts the cross-database combined query into a push-down joint query to be completed on the physical database.

Table Shard

The physical data table for each shard.

Transit Traffic

DRDS network traffic sent to the application result set.

Transmission Control Protocol

A connection-oriented and reliable transport layer communication protocol that is based on byte streams.

TCP

Transparent Data Encryption

Real-time I/O encryption and decryption of the entire database (including data and log files). It is completely transparent to applications and does not require any modifications to existing applications.

TDE

U

URL Forward

The action of automatically redirecting you to a preset address when you access a domain name. It includes explicit and implicit URL forwarding.

V

VIPServer

The service IP address mapping and environment management system. Based on centralized configurations, it provides users with routing information to achieve load balancing that does not depend on gateways. Multiple mapping policies (round robin, round robin of services deployed in a data center, and round robin of services deployed within a network segment) are supported. A health check mechanism automatically removes unhealthy hosts. The invocation between clusters is transparent. A certain amount of data, such as caller information and the number of invocation requests, is also provided.

Virtual Boarder Router (VBR)

The mapping of your leased line in VPCs. It can be regarded as a VRouter, or a forwarding bridge between an on-premises IDC and a VPC,

Virtual Extensible LAN (VXLAN)

A VXLAN-based Layer-2 network. A single VXLAN belongs to a large VXLAN network pool. VXLANs are isolated from each other at Layer-2.

VXLAN

Virtual Local Area Network

An emerging data exchange technology that logically divides Local Area Network (LAN) devices into segments to implement the virtual work group.

VLAN

Virtual Private Cloud

The custom private network created based on Alibaba Cloud. Full logical isolation is achieved between VPCs. You can create and manage cloud product instances, such as Elastic Compute Service (ECS), Relational Database Service (RDS), and Object Storage Service (OSS), in your created VPCs.

VPC

VRouter

A networking hub that uses Linux-specific VM instances to implement multiple network services.

VServer Group

A group of back-end servers that can be customized and managed in the listener dimension. They allow listeners under an SLB instance to distribute different requests to different back-end servers depending on the configured forwarding rule.

VSwitch

A basic network device of Virtual Private Cloud (VPC), which can connect to different cloud product instances and must be specified when you create a cloud product instance in VPC.

W

Weighted Least Connections

A scheduling algorithm that Server Load Balancer uses to forward traffic. In addition to the weight set to each back-end server, the actual loads (number of connections) of the back-end server must be considered. If the weights of the back-end servers are the same, a new request is sent to the server with the fewer connections.

WLC

Weighted Round Robin

A scheduling algorithm that Server Load Balancer uses to forward traffic. Back-end servers with higher weights receive more requests than those with lower weights.

WRR