# Alibaba Cloud

# Data Transmission Service Data Migration

Document Version: 20220712

C-J Alibaba Cloud

### Legal disclaimer

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

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized 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 because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should 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 this document based on the "status quo", "being defective", and "existing functions" of its products and services. 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 take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, 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 and/or its affiliates 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 directly contact Alibaba Cloud for any errors of this document.

## **Document conventions**

Style	Description	Example			
<u>↑</u> Danger	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.			
O Warning	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.			
C) Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.			
? Note	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 listinstanceid Instance_ID			
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]			
{} or {alb}	This format is used for a required value, where only one item can be selected.	switch {active stand}			

## Table of Contents

1.Data type mappings between heterogeneous databases	<mark>0</mark> 8
2.Precautions and limits	29
2.1. Precautions and limits for migrating data from a MySQL d	29
2.2. Precautions and limits for migrating data from a PolarDB	41
2.3. Precautions and limits for migrating data from a PolarDB	53
2.4. Limits for using PolarDB-X as the source instance	58
2.5. Solutions for using DRDS as the source instance	62
2.6. Precautions and limits for migrating data from an Oracle	65
2.7. Precautions and limits for migrating data from a PostgreS	77
2.8. Precautions and limits for migrating data from an SQL Se	90
2.9. Precautions and limits for migrating data from a Redis da	96
2.10. Precautions and limits for migrating data from a MongoD	99
2.11. Precautions and limits for migrating data from a Db2 for	103
2.12. Precautions and limits for migrating data from a Teradat	106
2.13. Precautions and limits for migrating data from an HBase	107
3.Prepare the database accounts for data migration	109
4. Migration task management	114
4.1. Object name mapping	114
4.2. Specify the capitalization of object names in the destinati	119
4.3. Filter the data to be migrated	121
4.4. View the progress of a data migration task	124
4.5. View the connection state and performance of full data m	126
4.6. View the connection status and performance of increment	128
4.7. Modify the transfer rate of full data migration	130
4.8. Fix a failed data migration task	131
4.9. Switch workloads to the destination database	133

4.10. Stop a data migration task	135
4.11. Data formats of a Kafka cluster	136
4.12. Specify the policy for migrating data to Kafka partitions	140
5.Precheck and troubleshooting	142
5.1. Source database connectivity	142
5.2. Destination database connectivity	143
5.3. Source database version	145
6.Migrate data from a self-managed database to Alibaba Cloud	146
6.1. Source database: MySQL	146
6.1.1. Migrate data from a self-managed MySQL database to	146
6.1.2. Migrate data from a self-managed MySQL database co	155
6.1.3. Migrate data from a self-managed MySQL database co	164
6.1.4. Migrate data from a self-managed MySQL database to	174
6.1.5. Migrate data from a self-managed MySQL database to	182
6.2. Source database: SQL Server	190
6.2.1. Migrate incremental data from a self-managed SQL Ser	190
6.2.2. Migrate full data from a self-managed SQL Server dat	200
6.3. Source database: Oracle	207
6.3.1. Migrate the schema of a self-managed Oracle database	207
6.3.2. Migrate data from a self-managed Oracle database to	211
6.3.3. Migrate data from a self-managed Oracle database to	218
6.3.4. Migrate data from a self-managed Oracle database to	226
6.3.5. Migrate data from a self-managed Oracle database to	235
6.3.6. Migrate data from a self-managed Oracle database to	243
6.3.7. Migrate data from a self-managed Oracle database to	251
6.3.8. Migrate data from a self-managed Oracle database to	259
6.4. Source database: PostgreSQL	270
6.4.1. Migrate incremental data from a self-managed Postgre	270

6.4.2. Migrate incremental data from a self-managed Postgre 280
6.4.3. Migrate full data from a self-managed PostgreSQL dat <sup>291</sup>
6.5. Source database: Redis 299
6.5.1. Migrate data from a self-managed Redis database to a 299
6.6. Source database: MongoDB 307
6.6.1. Migrate a self-managed standalone MongoDB database 307
6.6.2. Migrate a self-managed MongoDB database that uses 315
6.6.3. Migrate a self-managed MongoDB database that uses 322
6.7. Source database: TiDB 332
6.7.1. Migrate incremental data from a self-managed TiDB da
6.7.2. Migrate full data from a self-managed TiDB database t 342
6.8. Source database: Db2 350
6.8.1. Migrate data from a self-managed Db2 database to an 350
7. Migrate data from a third-party cloud to Alibaba Cloud 359
7.1. Migrate data from an Amazon RDS for MySQL instance to 359
7.2. Migrate data from an Amazon RDS for Oracle instance to 369
7.3. Migrate incremental data from an Amazon RDS for Postgr 381
7.4. Migrate full data from an Amazon RDS for PostgreSQL ins <sup>391</sup>
7.5. Migrate full data from an Amazon RDS for SQL Server ins400
7.6. Migrate data from an Amazon Aurora MySQL cluster to an408
7.7. Migrate data from an Amazon Aurora MySQL cluster to a 420
7.8. Migrate full data from an Amazon Aurora PostgreSQL inst 431
7.9. Migrate data from a MongoDB Atlas database to an Apsar 440
8.Migrate data between instances of the same Alibaba Cloud ac 448
8.1. Migrate data between RDS instances 448
8.2. Migrate data from an ApsaraDB RDS for MariaDB TX insta 459
8.3. Migrate data from an ApsaraDB RDS for PostgreSQL insta 468
8.4. Migrate data from an ApsaraDB RDS for MySQL instance 476

8.5. Migrate data from an ApsaraDB RDS for PPAS instance to...-484 8.6. Migrate data between PolarDB for MySQL clusters ------ 493 8.7. Migrate data from a PolarDB for MySQL cluster to an Aps... ------ 500 8.8. Migrate data between PolarDB for Oracle clusters ------ 508 8.9. Migrate data from a standalone instance to a replica set ... 515 8.10. Migrate data from an ApsaraDB for MongoDB replica set ........ 522 8.11. Migrate data between ApsaraDB for MongoDB instances a....... 528 8.12. Migrate data from a MaxCompute project to an ApsaraDB..----- 536 9. Migrate data between instances of different Alibaba Cloud acc... 543 9.1. Migrate data between PolarDB for MySQL clusters of differ... 543 9.2. Migrate data between ApsaraDB for MongoDB instances o... 552 9.3. Migrate data between RDS instances of different Alibaba ... ------ 560 10.Migrate data from Alibaba Cloud to a self-managed database 44444 569 10.1. Migrate data from an ApsaraDB RDS for MySQL instance ......... 569 10.2. Migrate data from an ApsaraDB RDS for MySQL instance... 577 10.3. Migrate data from a PolarDB for MySQL cluster to a self... 585 10.4. Migrate data from a PolarDB for PostgreSQL cluster to a... 593 11.Migrate data between self-managed databases ------ 600 11.1. Migrate data between self-managed Oracle databases ------ 600

# 1.Data type mappings between heterogeneous databases

Heterogeneous databases have different data types. When Data Transmission Service (DTS) migrates data between heterogeneous databases, DTS converts the data types of the source database to those of the destination database. This topic lists the data type mappings for you to evaluate the impact of data migration on your business.

#### Overview

You can view the data type mappings between heterogeneous databases based on the following migration scenarios:

- Migrate data from a PolarDB for MySQL cluster, an ApsaraDB RDS for MySQL instance, or a selfmanaged MySQL database
- Migrate data from an Oracle database
- Migrate data from an SQL Server database
- Migrate data from a self-managed TiDB database
- Migrate data from a Db2 for LUW database
- Migrate data from a Db2 for i database
- Migrate data from a Teradata database

## Migrate data from a PolarDB for MySQL cluster, an ApsaraDB RDS for MySQL instance, or a self-managed MySQL database

The following tables list the data type mappings between MySQL and AnalyticDB databases. The source instance can be a PolarDB for MySQL cluster, an ApsaraDB RDS for MySQL instance, or a self-managed MySQL database. The destination instance can be an AnalyticDB for MySQL cluster V2.0 or an AnalyticDB for PostgreSQL instance.

**Note** If the value range of the data to be migrated from the source instance exceeds the range supported by DTS, the accuracy of the data migrated to the destination instance decreases.

Туре	Data type of the source instance	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL	Data type of a Message Queue for Apache Kafka instance or a self-managed Kafka cluster
	BIT[(M)]	1 ~ 64	VARCHAR	BIT	
	TINYINT[(M)]	-128 ~ 127	TINYINT	SMALLINT	
	T INYINT [(M)] [UNSIGNED]	0 ~ 255	SMALLINT	SMALLINT	
	SMALLINT[(M)]	-32768 ~ 32767	SMALLINT	SMALLINT	

### Data Migration Data type mappings between heterogeneous databases

Туре	Data type of the source instance	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL	Data type of a Message Queue for Apache Kafka instance or a self-managed Kafka cluster
Integer	SMALLINT[(M)] [UNSIGNED]	0 ~ 65535	INT	INT EGER	
Integer	MEDIUMINT[(M) ]	-8388608 ~ 8388607	INT	INT EGER	
	MEDIUMINT[(M) ] [UNSIGNED]	0 ~ 16777215	INT	INT EGER	
	INT[(M)]	-2147483648 ~ 2147483647	INT	INT EGER	
	INT [(M)] [UNSIGNED]	0 ~ 4294967295	BIGINT	BIGINT	
	BIGINT [(M)]	- 922337203685 4775808 ~ 922337203685 4775807	BIGINT	BIGINT	
	BIGINT [(M)] [UNSIGNED]	0 ~ 184467440737 09551615.	DECIMAL(20,0)	NUMERIC(20)	
	DECIMAL[(M[,D] )]	M: 0 to 65 D: 0 to 30	DECIMAL[(M[,D] )]	DECIMAL	
	FLOAT (p)	1.175494351E- 38 ~ 3.402823466E +38	FLOAT	REAL	
	DOUBLE[(M,D)]	2.2250738585 072014E-308 ~ 1.7976931348 623157E+308	DOUBLE	DOUBLE PRECISION	
Decimal					

#### Dat a Transmission Service

Туре	Data type of the source instance	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL	Data type of a Message Queue for Apache Kafka instance or a self-managed Kafka cluster	
	DATE	1000-01- 01~9999-12- 31 <b>?</b> Note The format is YYYY- MM-DD, in UT C.	DATE	DATE		
	DAT ET IME[(fsp )]	1000-01-01 00:00:00.0000 00 ~ 9999-12- 31 23:59:59.9999 99 99 <b>??</b> Note The format is YYYY- MM-DD hh:mm:ss [.fraction], in UT C.	DAT ET IME	TIMESTAMP		
Date and time						

#### Data Migration Data type mappings between heterogeneous databases

Туре	Data type of the source instance	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL	Data type of a Message Queue for Apache Kafka instance or a Consistent self-managed with the data Kafka cluster types of
	T IMEST AMP[(fs p)]	1970-01-01 00:00:01.0000 00 ~ 2038-01- 19 03:14:07.9999 99 99 <b>?</b> <b>Note</b> The format is YYYY- MM-DD hh:mm:ss [.fraction], in UT C.	T IMEST AMP	T IMEST AMP WIT H T IME ZONE	MySQL or PolarDB for MySQL
	TIME[(fsp)]	- 838:59:59.000 000 ~ 838:59:59.000 000	TIME	TIME	
	YEAR[(4)] 1901 to 2155, or 0000		INT	INT EGER	
	CHAR[(M)]	HAR[(M)] 0 to 255 characters		CHAR	
	VARCHAR(M) 0 to 65,535 characters		VARCHAR	VARCHAR	
	BINARY[(M)]	0 to 255 bytes	VARBINARY	BYTEA	
	VARBINARY(M)	0 to 65,535 bytes	VARBINARY	BYTEA	

Туре	Data type of the source instance	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL	Data type of a Message Queue for Apache Kafka instance or a self-managed Kafka cluster
	TINYBLOB	255 (2^8 - 1) bytes	VARBINARY	BYTEA	
	TINYTEXT	255 (2^8 - 1) characters	VARCHAR	TEXT	
	BLOB	65,535 (2^16 - 1) bytes	VARBINARY	BYTEA	
String	TEXT	65,535 (2^16 - 1) characters	VARCHAR	TEXT	
	MEDIUMBLOB	16,777,215 (2^24 - 1) bytes	VARBINARY	BYTEA	
	MEDIUMT EXT	16,777,215 (2^24 - 1) characters	VARCHAR	TEXT	
	LONGBLOB	4,294,967,295 or 4 GB (2^32 - 1) bytes	VARBINARY	BYTEA	
	LONGT EXT	4,294,967,295 or 4 GB (2^32 - 1) characters	VARCHAR	TEXT	
	ENUM('value1', 'value2',)	An ENUM column can have a maximum of 65,535 distinct elements.	VARCHAR	VARCHAR(128)	
	SET('value1','v alue2',)	A SET column can have a maximum of 64 distinct elements.	VARCHAR	VARCHAR(128)	
	GEOMET RY	Geometry values of any type	VARBINARY	POLYGON	
	POINT	None	VARBINARY	POINT	

Туре	Data type of the source instance	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL	Data type of a Message Queue for Apache Kafka instance or a self-managed Kafka cluster
Spatial	LINEST RING	None	VARBINARY	PATH	
	POLYGON	None	VARBINARY	POLYGON	
	MULT IPOINT	None	VARBINARY	POLYGON	
	MULT ILINEST RI NG	None	VARBINARY	PATH	
	MULT IPOLY GO N	None	VARBINARY	POLYGON	
	GEOMET RYCOL LECT ION	A collection of geometry values of any type	VARBINARY	POLYGON	
JSON	JSON	None	JSON	JSON	

#### Migrate data from an Oracle database

The following table lists the data type mappings between Oracle and heterogeneous databases. The destination instance can be a self-managed MySQL database, a PolarDB for MySQL cluster, an AnalyticDB for PostgreSQL instance, or a PolarDB for Oracle cluster.

**Note** If the value range of the data to be migrated from the source instance exceeds the range supported by DTS, the accuracy of the data migrated to the destination instance will decrease.

Туре	Data type of Oracle	Value range	Data type of MySQL, PolarDB for MySQL, or PolarDB-X	Data type of ApsaraDB RDS for PPAS	Data type of AnalyticD B for MySQL	Data type of AnalyticD B for PostgreS QL	Data type of PolarDB for Oracle
------	------------------------	----------------	--	--	--	---	--

Туре	Data type of Oracle	Value range	Data type of MySQL, PolarDB for MySQL, or PolarDB-X	Data type of ApsaraDB RDS for PPAS	Data type of AnalyticD B for MySQL	Data type of AnalyticD B for PostgreS QL	Data type of PolarDB for Oracle
Numeric	NUMBER(p ,s)	1 to 22 bytes The argument p indicates the precision. Valid values: 1 to 38. The argument s indicates the scale. Valid values: - 84 to 127.	DECIMAL[( p[,s])]	NUMBER[( p[,s])]	DECIMAL   T INYINT   SMALLINT   INT EGER   BIGINT	DECIMAL   T INYINT   SMALLINT   INT EGER   BIGINT	NUMBER(p ,s)
	FLOAT (p)	1 to 22 bytes The variable p indicates a pointer. Valid values: 1 to 126 bits.	DOUBLE	DOUBLE PRECISION	DOUBLE	DOUBLE PRECISION	DOUBLE PRECISION
	BINARY_FL OAT	A 32-bit floating- point number (4 bytes)	DECIMAL(6 5,8)	REAL	DOUBLE	DOUBLE PRECISION	REAL
	BINARY_D OUBLE	A 64-bit floating- point number (8 bytes)	DOUBLE	DOUBLE PRECISION	DOUBLE	DOUBLE PRECISION	DOUBLE PRECISION
	DATE	None	DAT ET IME	DATE	DATETIME	T IMEST AM P(0)	DATE

#### Data Migration Data type mappings between heterogeneous databases

Туре	Data type of Oracle	Value range	Data type of MySQL, PolarDB for MySQL, or PolarDB-X	Data type of ApsaraDB RDS for PPAS	Data type of AnalyticD B for MySQL	Data type of AnalyticD B for PostgreS QL	Data type of PolarDB for Oracle
Date and time	TIMESTAM P [(fraction al_second s_precisio n)]	None	DAT ET IME[ (fractiona l_seconds _precision )]	T IMEST AM P [(fraction al_second s_precisio n)]	DAT ET IME	T IMEST AM P	TIMEST AM P [(fraction al_second s_precisio n)]
	T IMEST AM P [(fraction al_second s_precisio n)] WIT H T IME ZONE	None	DAT ET IME[ (fractiona l_seconds _precision )]	TIMESTAM P [(fraction al_second s_precisio n)] WIT H TIME ZONE	T IMEST AM P	TIMESTAM P WITH TIME ZONE	TIMESTAM P [(fraction al_second s_precisio n)] WITH TIME ZONE
	TIMESTAM P [(fraction al_second s_precisio n)] WITH LOCAL TIME ZONE	None	DAT ET IME[ (fractiona l_seconds _precision )]	TIMESTAM P [(fraction al_second s_precisio n)] WITH TIME ZONE	DAT ET IME	TIMESTAM P WITH TIME ZONE	TIMEST AM P [(fraction al_second s_precisio n)] WIT H TIME ZONE
	INT ERVAL YEAR [(year_pre cision)] TO MONTH	None	Not supporte d	Not supporte d	VARCHAR	VARCHAR( 32)	INTERVAL
	INT ERVAL DAY [(day_pre cision)] TO SECOND [(fraction al_second s_precisio n)]	None	Not supporte d	Not supporte d	VARCHAR	VARCHAR( 32)	INTERVAL
	CHAR [(size [BYTE  CHAR])]	2,000 bytes	CHAR[(n)]	CHAR[(n)]	VARCHAR	CHAR	CHAR [(size [BYTE  CHAR])]

#### Dat a Transmission Service

Туре	Data type of Oracle	Value range	Data type of MySQL, PolarDB for MySQL, or PolarDB-X	Data type of ApsaraDB RDS for PPAS	Data type of AnalyticD B for MySQL	Data type of AnalyticD B for PostgreS QL	Data type of PolarDB for Oracle
	NCHAR[(si ze)]	2,000 bytes	NAT IONAL CHAR[(n)]	NCHAR[(n) ]	VARCHAR	VARCHAR	NCHAR[(si ze)]
	VARCHAR2 (size [BYTE  CHAR])	If MAX_STRI NG_SIZE is set to EXTENDED , the maximum size is 32,767 bytes. If MAX_STRI NG_SIZE is set to STANDAR D, the maximum size is 4,000 bytes.	VARCHAR( n)	VARCHAR 2[(n)]	VARCHAR	VARCHAR	VARCHAR 2(size [BYTE   CHAR])
	NVARCHA R2(size)	If MAX_STRI NG_SIZE is set to EXTENDED , the maximum size is 32,767 bytes. If MAX_STRI NG_SIZE is set to STANDAR D, the maximum size is 4,000 bytes.	NAT IONAL VARCHAR[ (n)]	VARCHAR 2[(n)]	VARCHAR	VARCHAR	NVARCHA R2(size)

#### Data Migration Data type mappings between heterogeneous databases

String Type	Data type of Oracle	Value range	Data type of MySQL, PolarDB for MySQL, or PolarDB-X	Data type of ApsaraDB RDS for PPAS	Data type of AnalyticD B for MySQL	Data type of AnalyticD B for PostgreS QL	Data type of PolarDB for Oracle
	LONG	The maximum size is 2 GB (2^31 - 1).	LONGT EX T	LONG	VARCHAR	TEXT	LONG
	RAW(size)	The maximum size is 32,767 bytes or 2,000 bytes.	VARBINAR Y(2000)	RAW(size)	VARBINAR Y	BYTEA	RAW(size)
	LONG RAW	The maximum size is 2 GB.	LONGBLO B	LONG RAW	VARBINAR Y	BYTEA	LONG RAW
	CLOB	The maximum size is (4 GB - 1) × DB_BLOCK _SIZE.	LONGT EX T	CLOB	VARCHAR	TEXT	CLOB
	NCLOB	The maximum size is (4 GB - 1) × DB_BLOCK _SIZE.	LONGT EX T	NCLOB	VARCHAR	TEXT	CLOB
	BLOB	The maximum size is (4 GB - 1) × DB_BLOCK _SIZE.	LONGBLO B	BLOB	VARBINAR Y	BYTEA	BLOB
	BFILE	The maximum size is 4 GB.	Not supporte d	Not supporte d	Not supporte d	Not supporte d	Not supporte d

Туре	Data type of Oracle	Value range	Data type of MySQL, PolarDB for MySQL, or PolarDB-X	Data type of ApsaraDB RDS for PPAS	Data type of AnalyticD B for MySQL	Data type of AnalyticD B for PostgreS QL	Data type of PolarDB for Oracle
JSON	JSON	The maximum size is 32 MB.	Not supporte d	Not supporte d	JSON	JSON	JSON
ROWID	ROWID	64 character s	Not supporte d	Not supporte d	ROWID	OID	VARCHAR
Spatial	Customiz ation required	DTS does not support					

#### ? Note

- Destination instance: MySQL, PolarDB for MySQL, or PolarDB-X
  - If a CHAR field in the source database is greater than 255 bytes in length, DTS converts this field to the VARCHAR(n) type in the destination database.
  - MySQL does not support the following Oracle data types: BFILE, INTERVAL YEAR TO MONTH, and INTERVAL DAY TO SECOND. Therefore, DTS does not convert these data types during schema migration.

If a table to be migrated contains these data types, schema migration fails. You must make sure that the columns with these data types are excluded from the objects to be migrated.

- The TIMESTAMP data type of MySQL databases does not contain the time zone information. However, the IMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE data types in Oracle databases provide the time zone information. Therefore, DTS converts the values of these data types to UTC time in the destination database.
- Destination instance: ApsaraDB RDS for PPAS ApsaraDB RDS for PPAS does not support the TIMESTAMP[(fractional\_seconds\_precision)] WITH LOCAL TIME ZONE data type. DTS converts the data of this type to UTC time and then stores the data in the destination ApsaraDB RDS for PPAS instance by using the TIMESTAMP[(fractional\_seconds\_precision)] WITH TIME ZONE data type.
- Destination instance: AnalyticDB for PostgreSQL If an Oracle data type is not supported by AnalyticDB for PostgreSQL, DTS converts the data type to BYTEA. If the conversion fails, DTS sets the field value to NULL.

#### Migrate data from an SQL Server database

The following table lists the data type mappings between SQL Server and AnalyticDB databases. The source instance can be a self-managed SQL Server database or an ApsaraDB RDS for SQL Server instance. The destination instance can be an AnalyticDB for MySQL cluster or an AnalyticDB for

#### Post greSQL instance.

**Note** If the value range of the data to be migrated from the source instance exceeds the range supported by DTS, the accuracy of the data migrated to the destination instance will decrease.

Туре	Data type of SQL Server	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL
	BIT	An INTEGER data type that can take a value of 1, 0, or NULL	BOOLEAN	BIT(1)
	TINYINT	0 to 255	TINYINT	SMALLINT
	SMALLINT	-32,768 (-2^15) to 32,767 (2^15 - 1)	SMALLINT	SMALLINT
Integer	INT	-2,147,483,648 (- 2^31) to 2,147,483,647 (2^31 - 1)	INT EGER	INT EGER
	BIGINT	- 9,223,372,036,854, 775,808 (-2^63) to 9,223,372,036,854, 775,807 (2^63 - 1)	BIGINT	BIGINT
	NUMERIC[ (p[ ,s] )]	-10^38 + 1 to 10^38 - 1 (1 <= p <= 38)	DECIMAL	DECIMAL
	DECIMAL[ (p[ ,s] )]	-10^38 + 1 to 10^38 - 1 (1 <= p <= 38)	DECIMAL	DECIMAL
Decimal	FLOAT	-1.79E + 308 to - 2.23E - 308, 0, and 2.23E - 308 to 1.79E + 308	DOUBLE	DOUBLE PRECISION
	REAL	-3.40E + 38 to - 1.18E - 38, 0, and 1.18E - 38 to 3.40E + 38	FLOAT	REAL
	MONEY	- 922,337,203,685,4 77.5808 to 922,337,203,685,4 77.5807	DECIMAL(19, 4)	DECIMAL(19, 4)

#### Dat a Transmission Service

Monetary Type	Data type of SQL Server	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL
	SMALLMONEY	-214,748.3648 to 214,748.3647	DECIMAL(10, 4)	DECIMAL(10, 4)
	DATE	0001-01-01 to 9999-12-31	DATE	DATE
Date and time	DAT ET IME	Date range: January 1, 1753 to December 31, 9999 Time range: 00:00:00 to 23:59:59.997	DAT ET IME	TIMESTAMP(3) WITHOUT TIME ZONE
	DAT ET IME2[ (fractional seconds precision) ]	Date range: January 1,1 CE to December 31, 9999 CE Time range: 00:00:00 to 23:59:59.99999999	DAT ET IME	TIMESTAMP(7) WITHOUT TIME ZONE
	DAT ET IMEOFFSET [ (fractional seconds precision) ]	Date range: January 1, 1 CE to December 31, 9999 CE Time range: 00:00:00 to 23:59:59.9999999 Time zone offset range: -14:00 to +14:00	TIMESTAMP	T IMEST AMP(7) WIT H T IME ZONE
	SMALLDAT ET IME	The time is based on a 24-hour day, with seconds always zero (:00) and without fractional seconds.	DAT ET IME	TIMESTAMP WITHOUT TIME ZONE
	TIME [ (fractional second scale) ]	00:00:00.0000000 to 23:59:59.99999999	TIME	T IME(7) WIT H T IME ZONE
	BINARY [ ( n ) ]	Valid values of n: 1 to 8000.	VARBINARY	BYTEA

#### Data Migration Data type mappings between heterogeneous databases

Туре	Data type of SQL Server	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL
	VARBINARY [ ( n   max) ]	Valid values of n: 1 to 8000. max indicates that the maximum storage size is 2^31 - 1 bytes.	VARBINARY	BYTEA
	CHAR [ ( n ) ]	Valid values of n: 1 to 8000. The storage size is n bytes.	VARCHAR	CHARACTER
String	VARCHAR [ ( n   max ) ]	Valid values of n: 1 to 8000. max indicates that the maximum storage size is 2^31 - 1 bytes (2 GB).	VARCHAR	CHARACTER
	NCHAR [ ( n ) ]	n defines the string size in byte- pairs. Valid values of n: 1 to 4000. The storage size is two times n bytes.	VARCHAR	CHARACT ER VARYING
	NVARCHAR [ ( n   max ) ]	n defines the string size in byte- pairs. Valid values of n: 1 to 4000. max indicates that the maximum storage size is 2^30 - 1 characters (2 GB).	VARCHAR	CHARACT ER VARYING
	NTEXT	Variable-length Unicode data with a maximum string length of 1,073,741,823 (2^30 - 1) bytes.	VARCHAR	TEXT
	TEXT	The maximum string length is 2,147,483,647 (2^31 - 1) bytes.	VARCHAR	TEXT

Туре	Data type of SQL Server	Value range	Data type of AnalyticDB for MySQL	Data type of AnalyticDB for PostgreSQL
	IMAGE	Variable-length binary data from 0 to 2,147,483,647 (2^31 - 1) bytes.	VARBINARY	BYTEA
Spatial (geography and geometry)	GEOGRAPHY	None	VARCHAR	Not supported
	GEOMET RY	None	VARCHAR	Not supported
XML	XML ( [ CONTENT   DOCUMENT ] xml_schema_colle ction )	None	VARCHAR	XML
	UNIQUEIDENT IFIER	None	VARCHAR	CHARACTER(36)
Other types	SQL_VARIANT	None	Not supported	Not supported
	HIERARCHYID	None	Not supported	Not supported
	SYSNAME	None	VARCHAR	CHARACTER VARYING(128)

#### Migrate data from a self-managed TiDB database

The following table lists the data type mappings between a self-managed TiDB database and a MySQL database.

Data type of TiDB	Data type of MySQL
BIGINT	BIGINT
BINARY	BINARY
BIT	BIT
BOOL\ BOOLEAN	TINYINT
CHAR	CHAR
DATE	DATE
DATETIME	DAT ET IME
DECIMAL	DECIMAL
DOUBLE	DOUBLE
ENUM	ENUM

Data type of TiDB	Data type of MySQL
FLOAT	FLOAT
INT	INT
INTEGER	INTEGER
JSON	JSON
MEDIUMBLOB/LONGBLOB T INYBLOB/BLOB/	MEDIUMBLOB/LONGBLOB T INYBLOB/BLOB/
MEDIUMINT	MEDIUMINT
SET	SET
SMALLINT	SMALLINT
T EXT / LONGT EXT	TEXT/LONGTEXT
ТІМЕ	TIME
TIMESTAMP	TIMESTAMP
TINYINT	TINYINT
VARBINARY	VARBINARY
VARCHAR	VARCHAR
YEAR	YEAR

#### Migrate data from a Db2 for LUW database

The following table lists the data type mappings between a Db2 for LUW database and a MySQL database.

**Note** If the value range of the data to be migrated from the source instance exceeds the range supported by DTS, the accuracy of the data migrated to the destination instance will decrease.

Туре	Data type of Db2 for LUW	Value range	Data type of MySQL
	SMALLINT	-32,768~+32,767	SMALLINT
	INTEGER	- 2,147,483,648~+2,147,4 83,647	INT
Integer			

#### Dat a Transmission Service

Туре	Data type of Db2 for Value range		Data type of MySQL
	BIGINT	- 9,223,372,036,854,775,8 08~ +9,223,372,036,854,775, 807	BIGINT
	DECIMAL(precision- integer, scale-integer)	p<=38	DECIMAL
Decimal	FLOAT (integer)	The value range is 1 to 53. If the integer is between 1 and 24 inclusive, the format is single precision floating-point. If the integer is between 25 and 53 inclusive, the format is double precision floating-point.	FLOAT
	DECFLOAT (precision- integer)	None	DECIMAL(65,10)
	DATE	0001-01-01~9999-12-31	DATE
	TIME	00:00:00~24:00:00	TIME
Date and time	T IMEST AMP(integer)	0001-01-01- 00.00.00.00000000000 ~9999-12-31- 24.00.00.000000000000; 0<=p<= 12	DATETIME
	CHARACT ER(integer)	254	CHAR   VARCHAR
	VARCHAR(integer)	32,672	VARCHAR
	CHARACTER(integer) FOR BIT DATA	254	BLOB
String	CLOB	2,147,483,647	LONGTEXT
	GRAPHIC (integer)	127	CHAR(length*4)
	VARGRAPHIC (integer)	16,336	CHAR(length*4)
	DBCLOB (integer)	1,073,741,823	VARCHAR   LONGT EXT
	BLOB	2,147,483,647	LONGBLOB
Other types	XML	2,147,483,647	VARCHAR   LONGT EXT

#### Migrate data from a Db2 for i database

#### The following table lists the data type mappings between a Db2 for i database and a MySQL database.

**Note** If the value range of the data to be migrated from the source instance exceeds the range supported by DTS, the accuracy of the data migrated to the destination instance will decrease.

Туре	Data type of Db2 for i	Value range	Data type of MySQL
	SMALLINT	-32,768~+32,767	SMALLINT
Integer	INTEGER	- 2,147,483,648~+2,147,4 83,647	INT
Integer	BIGINT	- 9,223,372,036,854,775,8 08~ +9,223,372,036,854,775, 807	BIGINT
	DECIMAL(precision- integer, scale-integer)	p<=63	DECIMAL
Docimal	NUMERIC	None	DECIMAL
Decimat	FLOAT (integer)	None	FLOAT
	DECFLOAT (precision- integer)	None	DECIMAL(65,10)
	DATE	0001-01-0~9999-12-31	DATE
	TIME	00:00:00~24:00:00	TIME
Date and time	TIMESTAMP(integer)	0001-01-01- 00.00.00.000000000000 to 9999-12-31- 24.00.00.000000000000 (0 <= p <= 12)	DATETIME
	CHAR(integer)	32,765	CHAR   VARCHAR
	VARCHAR(integer)	32,739	VARCHAR
	CHAR(integer) FOR BIT DATA	None	BLOB
	CLOB	2,147,483,647	LONGTEXT
	GRAPHIC (integer)	16,382	CHAR
String	VARGRAPHIC (integer)	16,369	VARCHAR

Туре	Data type of Db2 for i	Value range	Data type of MySQL
	DBCLOB (integer)	1,073,741,823	LONGT EXT
	BINARY	32,765	BINARY
	VARBIN	32,739	VARBINARY
	BLOB	2,147,483,647	LONGBLOB
	DATALINK	None	VARCHAR   LONGT EXT
Other types	ROWID	40	VARCHAR   LONGT EXT
	XML	2,147,483,647	VARCHAR   LONGT EXT

#### Migrate data from a Teradata database

The following table lists the data type mappings between a Teradata database and an AnalyticDB for PostgreSQL instance.

Data type of Teradata	Data type of AnalyticDB for PostgreSQL
BYTEINT	SMALLINT
SMALLINT	SMALLINT
BIGINT	BIGINT
INTEGER	INT EGER
DATE	DATE
JSON	JSON
XML	XML
CLOB	text
Float	real
CHAR	CHAR
VARCHAR	VARCHAR
Timestamp	Timestamp
TIME	TIME
Timestamp With Time Zone	Timestamp With Time Zone
Time With Time Zone	Time With Time Zone
Decimal	Decimal

Data type of Teradata	Data type of AnalyticDB for PostgreSQL
Number	numeric
BYTE	bytea
VARBYTE	bytea
BLOB	bytea
PERIOD	varchar(100)
INTERVAL	varchar(100)

Data type of Teradata	Data type of AnalyticDB for PostgreSQL
SMALLINT	SMALLINT
INTEGER	INT
BIGINT	BIGINT
DECIMAL(precision-integer, scale-integer)	DECIMAL
NUMERIC	DECIMAL
FLOAT (integer)	FLOAT
DECFLOAT (precision-integer)	DECIMAL(65,10)
DATE	DATE
TIME	TIME
TIMESTAMP(integer)	DATETIME
CHAR(integer)	CHAR   VARCHAR
VARCHAR(integer)	VARCHAR
CHAR(integer) FOR BIT DATA	BLOB
CLOB	LONGTEXT
GRAPHIC (integer)	CHAR
VARGRAPHIC (integer)	VARCHAR
DBCLOB (integer)	LONGTEXT
BINARY	BINARY
VARBIN	VARBINARY

Data type of Teradata	Data type of AnalyticDB for PostgreSQL
BLOB	LONGBLOB
DATALINK	VARCHAR   LONGT EXT
ROWID	VARCHAR   LONGT EXT
XML	VARCHAR   LONGT EXT

## 2.Precautions and limits 2.1. Precautions and limits for migrating data from a MySQL database

This topic describes the precautions and limits when you migrate data from a MySQL database, such as a self-managed MySQL database and an ApsaraDB RDS for MySQL instance. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from a MySQL database

Take note of precautions and limits in the following data migration scenarios:

- Migrate data between MySQL databases
- Migrate data from a MySQL database to a PolarDB for MySQL cluster
- Migrate data from a MySQL database to a PolarDB-X instance
- Migrate data from a MySQL database to an AnalyticDB for MySQL cluster
- Migrate data from a MySQL database to a self-managed Kafka cluster
- Migrate data from a MySQL database to a DataHub project

#### Migrate data between MySQL databases

The following table describes the precautions and limits when you migrate data between MySQL databases, such as self-managed MySQL databases and ApsaraDB RDS for MySQL instances.

Category De

Description

Category	Description
Limits on the source database	• The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.
	• The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
	<ul> <li>If you select tables as the objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
	<ul> <li>The binary logging feature is enabled. The value of the binlog_format parameter is set to row. The value of the binlog_row_image parameter is set to full. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>
	Notice If the self-managed MySQL database is deployed in a dual-primary cluster, you must set log_slave_updates to ON. This ensures that Data Transmission Service (DTS) can obtain all binary logs.
	<ul> <li>For an incremental data migration task, binary logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, binary logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>

Category	Description
	• To ensure compatibility, we recommend that you use the same engine versions for the source and destination MySQL databases.
	• Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.
Other	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.
Other limits	• You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meet your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.
	<ul> <li>If the source database is a self-managed MySQL database, take note of the following limits:</li> </ul>
	<ul> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> </ul>
Special cases	<ul> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no DML operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the synchronization task is too high, you can perform a DML operation on the source database to update the latency.</li> </ul>
	<b>Note</b> If you select an entire database as the object to migrate, you can create a heartbeat table. The heartbeat table is updated or receives data every second.
	<ul> <li>If the destination database runs on an ApsaraDB RDS for MySQL instance, take note of the following limits:</li> <li>DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance.</li> <li>However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task. For more information, see Create a database on an ApsaraDB RDS for MySQL instance.</li> </ul>

#### Migrate data from a MySQL database to a PolarDB for MySQL cluster

The following table describes the precautions and limits when you migrate data from a MySQL database to a PolarDB for MySQL cluster.

Category	Description
Limits on the source database	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
	<ul> <li>The binary logging feature is enabled. The value of the binlog_format parameter is set to row. The value of the binlog_row_image parameter is set to full. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> <li>Notice If the self-managed MySQL database is deployed in a dual-primary cluster, you must set log_slave_updates to ON. This ensures that Data Transmission Service (DTS) can obtain all binary logs.</li> </ul>
	<ul> <li>For an incremental data migration task, binary logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, binary logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	<ul> <li>Limits on operations:</li> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>

Category	Description
Ot her limits	<ul> <li>To ensure compatibility, we recommend that you use the same engine versions for the source and destination MySQL databases.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meet your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the</li> </ul>
	accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.
Special cases	<ul> <li>If the source database is a self-managed MySQL database, take note of the following limits:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no DML operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the synchronization task is too high, you can perform a DML operation on the source database to update the latency.</li> </ul>
	<b>Note</b> If you select an entire database as the object to migrate, you can create a heartbeat table. The heartbeat table is updated or receives data every second.
	<ul> <li>If the destination database runs on a PolarDB for MySQL cluster, take note of the following limits:</li> <li>DTS automatically creates a destination database in the PolarDB for MySQL cluster.</li> <li>However, if the name of the source database is invalid, you must manually create a database in the PolarDB for MySQL cluster before you configure the data migration task.</li> <li>For more information, see Database Management.</li> </ul>

#### Migrate data from a MySQL database to a PolarDB-X instance

The following table describes the precautions and limits.

Category

Description

Category	Description
	• The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.
	• The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
	• If you select tables as the objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.
	• If you want to migrate incremental data, you must make sure that the following requirements are met:
	<ul> <li>The binary logging feature is enabled. The value of the binlog_format parameter is set to row. The value of the binlog_row_image parameter is set to full. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>
Limits on the source database	<b>Notice</b> If the self-managed MySQL database is deployed in a dual- primary cluster, you must set log_slave_updates to ON. This ensures that Data Transmission Service (DTS) can obtain all binary logs.
	<ul> <li>For an incremental data migration task, binary logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, binary logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select full data migration and incremental data migration as the migration types.</li> </ul>

Category	Description
Other limits	<ul> <li>The storage type of the PolarDB-X instance must be ApsaraDB RDS for MySQL (private custom RDS instance). PolarDB for MySQL cannot be used as the storage type.</li> <li>Before you configure a data migration task, you must create databases and tables in the destination instance.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meet your business requirements. DTS uses the ROUND (COLUM N, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.</li> </ul>
Special cases	<ul> <li>If the source database is a self-managed MySQL database, take note of the following limits:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no DML operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the synchronization task is too high, you can perform a DML operation on the source database to update the latency.</li> <li>Note If you select an entire database as the object to migrate, you can create a heartbeat table. The heartbeat table is updated or receives data every second.</li> </ul>

# Migrate data from a MySQL database to an AnalyticDB for MySQL cluster

The following table describes the precautions and limits.

Category	Description
Limits on the source database	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> <li>The binary logging feature is enabled. The value of the binlog format parameter is set</li> </ul>
	to row. The value of the binlog_row_image parameter is set to full. Otherwise, error messages are returned during precheck and the data migration task cannot be started.  Notice If the self-managed MySQL database is deployed in a dual-primary cluster, you must set log_slave_updates to ON. This ensures that Data Transmission Service (DTS) can obtain all binary logs.
	<ul> <li>For an incremental data migration task, binary logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, binary logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	<ul> <li>Limits on operations:</li> <li>During schema migration and full data migration, do not perform DDL operations to</li> </ul>
	<ul> <li>change the schemas of databases or tables. Otherwise, the data migration task fails.</li> <li>During data migration, do not perform DDL operations to add comments. Example: AL TER TABLE table_name COMMENT='Table comment';</li> <li>Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>
Category	Description
------------------	---
Other limits	<ul> <li>Prefix indexes cannot be migrated. If the source database contains prefix indexes, data may fail to be migrated.</li> <li>Due to the limits of , if the disk space usage of the nodes in an cluster reaches 80%, the task is delayed and error messages are returned. We recommend that you estimate the required disk space based on the objects to migrate. You must make sure that the destination cluster has sufficient storage space.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meet your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.</li> </ul>
Special cases	<ul> <li>If the source database is a self-managed MySQL database, take note of the following limits:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no DML operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the synchronization task is too high, you can perform a DML operation on the source database to update the latency.</li> <li>Note If you select an entire database as the object to migrate, you can create a heartbeat table. The heartbeat table is updated or receives data every second.</li> </ul>

# Migrate data from a MySQL database to a self-managed Kafka cluster

The following table describes the precautions and limits.

Category

Category	Description
	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> <li>The binary logging feature is enabled in the ApsaraDB RDS console. For more information, see Modify the parameters of an ApsaraDB RDS for MySQL instance. The value of the binlog_format parameter is set to row. The value of the binlog_row_image parameter is set to full. Otherwise, error messages are returned during precheck and the data synchronization task cannot be started.</li> </ul>
Limits on the source database	<ul> <li>Notice If the self-managed MySQL database is deployed in a dual-primary cluster, you must set log_slave_updates to ON. This ensures that DTS can obtain all binary logs. For more information, see Create an account for a user-created MySQL database and configure binary logging.</li> <li>For an incremental data migration, binary logs of the source database must be stored for more than 24 hours. For a schema and incremental data migration, binary logs of the source database must be stored for at least seven days. After schema migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of binary logs in accordance with the preceding requirements. Otherwise, the SLA of DTS does not ensure service reliability and performance. For more information about binary log files and log backup files of an ApsaraDB RDS for MySQL instance, see Introduction to binary log files and log backup files of an ApsaraDB RDS for MySQL instance, the data migration task fails.</li> <li>It you perform only full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration types.</li> </ul>

Category	Description
Limits	<ul> <li>Before you configure a data migration task, you must create a Kafka cluster and the Kafka version must be 0.10.1.0 to 2.0.</li> <li>You can select only tables as objects to migrate.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meet your business requirements. DTS uses the ROUND (COLUM N, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.</li> <li>We recommend that you do not use tools other than DTS to write data to the destination database. Otherwise, data inconsistency between the source and destination database.</li> </ul>
Special cases	<ul> <li>If the source database is a self-managed MySQL database, take note of the following limits:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no DML operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the synchronization task is too high, you can perform a DML operation on the source database to update the latency.</li> <li>Note If you select an entire database as the object to migrate, you can create a heartbeat table. The heartbeat table is updated or receives data every second.</li> </ul>

#### Migrate data from a MySQL database to a DataHub project

The following table describes the precautions and limits.

Category	Description
Limits on the source database	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> <li>The binary logging feature is enabled in the ApsaraDB RDS for MySQL instance. The value of the binlog_format parameter is set to row. The value of the binlog_row_image parameter is set to full. Otherwise, error messages are returned during precheck and the data synchronization task cannot be started.</li> </ul>
	<b>Notice</b> If the self-managed MySQL database is deployed in a dual-primary cluster, you must set log_slave_updates to ON. This ensures that DTS can obtain all binary logs. For more information, see <b>Create an account for a user-created MySQL database and configure binary logging</b> .
	<ul> <li>For an incremental data migration, binary logs of the source database must be stored for more than 24 hours. For a schema and incremental data migration, binary logs of the source database must be stored for at least seven days. After schema migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of binary logs in accordance with the preceding requirements. Otherwise, the SLA of DTS does not ensure service reliability and performance. For more information about binary log files and log backup files of an ApsaraDB RDS for MySQL instance, see Introduction to binary log files and log backup files of an ApsaraDB RDS for MySQL instance.</li> </ul>
	• Limits on operations: During schema migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.

Category	Description
Other limits	<ul> <li>Only tables can be selected as objects to migrate.</li> <li>We recommend that you do not use gh-ost or pt-online-schema-change to perform DDL operations on objects during data migration. Otherwise, data migration may fail.</li> <li>If you use only DTS to write data to the destination database, you can use Data Management (DMS) to perform online DDL operations on source tables during data synchronization. For more information, see Change schemas without locking tables.</li> </ul>
	<b>Warning</b> If you use tools other than DTS to write data to the destination database, we recommend that you do not use DMS to perform online DDL operations. Otherwise, data loss may occur in the destination database.
	<ul> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meet your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.</li> </ul>
Special cases	<ul> <li>If the source database is a self-managed MySQL database, take note of the following limits:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no DML operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the synchronization task is too high, you can perform a DML operation on the source database to update the latency.</li> </ul>
	<b>Note</b> If you select an entire database as the object to migrate, you can create a heartbeat table. The heartbeat table is updated or receives data every second.

## 2.2. Precautions and limits for migrating data from a PolarDB for MySQL cluster

This topic describes the precautions and limits when you migrate data from a cluster. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from a PolarDB for MySQL cluster

You can view the precautions and limits based on the following migration scenarios:

- Migrate data between PolarDB for MySQL clusters
- Migrate data from a PolarDB for MySQL cluster to an ApsaraDB RDS for MySQL instance or a selfmanaged MySQL database
- Migrate data from a PolarDB for MySQL cluster to a PolarDB-X instance
- Migrate data from a PolarDB for MySQL cluster to an AnalyticDB for MySQL cluster
- Migrate data from a PolarDB for MySQL cluster to a self-managed Oracle database
- PolarDB MySQL迁移至DataHub

#### Migrate data between PolarDB for MySQL clusters

The following table describes the precautions and limits.

Category

Category	Description
	• Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.
	• The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
	• If you select tables as the objects to be migrated and you need to edit tables (such as rename tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.
	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
	<ul> <li>The binary logging feature must be enabled. The value of the binlog_format parameter must be set to row. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>
Limits on the source database	<ul> <li>If you perform only incremental data migration, the binary logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the binary logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>

Category	Description
Other limits	<ul> <li>To ensure compatibility, we recommend that you use the same MySQL version for the source and destination clusters.</li> <li>Read-only nodes of the source cluster cannot be migrated.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (C OLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 38 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.</li> </ul>

# Migrate data from a PolarDB for MySQL cluster to an ApsaraDB RDS for MySQL instance or a self-managed MySQL database

The following table describes the precautions and limits.

Category

Category	Description
	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as rename tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
	<ul> <li>The binary logging feature must be enabled. The value of the binlog_format parameter must be set to row. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>
Limits on the source database	<ul> <li>If you perform only incremental data migration, the binary logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the binary logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>
	Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>

Category	Description
Precautions	<ul> <li>Read-only nodes of the source cluster cannot be migrated.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLU MN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits</li> </ul>
	<ul> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.</li> </ul>
Special cases	If you migrate data to an ApsaraDB RDS for MySQL instance, DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task. For more information, see Create a database on an ApsaraDB RDS for MySQL instance.

# Migrate data from a PolarDB for MySQL cluster to a PolarDB-X instance

The following table describes the precautions and limits.

Category

Category	Description
	• Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.
	• The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
	• If you select tables as the objects to be migrated and you need to edit tables (such as rename tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.
	• If you want to migrate incremental data, you must make sure that the following requirements are met:
	<ul> <li>The binary logging feature must be enabled. The value of the binlog_format parameter must be set to row. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>
Limits on the source database	<ul> <li>If you perform only incremental data migration, the binary logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the binary logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Full Data Migration and Incremental Data Migration as the migration types.</li> </ul>
	<ul> <li>Incremental DDL operations cannot be migrated. If you perform DDL operations in the source database during incremental data migration, the data migration task fails. If you need to perform DDL operations, we recommend that you perform the operations in the destination database and then perform the operations in the source database.</li> </ul>

Category	Description
	<ul> <li>Schema migration is not supported. Before you configure a data migration task, you must create databases and tables in the destination instance.</li> <li>The storage type of the PolarDB-X instance must be ApsaraDB RDS for MySQL (private custom RDS instance). cannot be used as the storage type.</li> <li>Read-only nodes of the source cluster cannot be migrated.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers</li> </ul>
Precautions	<ul> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the tablespace of the destination database is larger than that of the source database.</li> </ul>
	• You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLU MN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

# Migrate data from a PolarDB for MySQL cluster to an AnalyticDB for MySQL cluster

The following table describes the precautions and limits.

Category

Category	Description
	• Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.
	<ul> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> </ul>
	<ul> <li>If you select tables as the objects to be migrated and you need to edit tables (such as rename tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	• If you want to migrate incremental data, you must make sure that the following requirements are met:
	<ul> <li>The binary logging feature must be enabled. The value of the binlog_format parameter must be set to row. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>
Limits on the source database	<ul> <li>If you perform only incremental data migration, the binary logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the binary logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>
	Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>During data migration, do not perform DDL operations to add comments, for example, ALTER TABLE table_name COMMENT='Table comment';</li> <li>Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>

Category	Description	
	• Prefix indexes cannot be migrated. If the source database contains prefix indexes, data may fail to be migrated.	
	<ul> <li>Read-only nodes of the source cluster cannot be migrated.</li> </ul>	
	• Due to the limits of , if the disk space usage of the nodes in an cluster reaches 80%, the task is delayed and error messages are returned. We recommend that you estimate the required disk space based on the objects that you want to migrate. You must make sure that the destination cluster has sufficient storage space.	
	• Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.	
Precautions	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the tablespace of the destination database is larger than that of the source database.	
	• You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLU MN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.	
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.	

# Migrate data from a PolarDB for MySQL cluster to a self-managed Oracle database

The following table describes the precautions and limits.

Category

Category	Description		
	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as rename tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrate the entire database.</li> </ul>		
	• If you want to migrate incremental data, you must make sure that the following requirements are met:		
	<ul> <li>The binary logging feature must be enabled. The value of the binlog_format parameter must be set to row. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>		
Limits on the source database	<ul> <li>If you perform only incremental data migration, the binary logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the binary logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the binary logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of binary logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>		
	• Limits on operations:		
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>		
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>		

Category	Description		
Precautions	<ul> <li>Read-only nodes of the source cluster cannot be migrated.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLU MN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.</li> </ul>		
Special cases	If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Connect an on-premises database to Alibaba Cloud and Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.		

#### PolarDB MySQL迁移至DataHub

具体注意事项及限制如下:

类型	说明		
----	----	--	--

类型	说明
源库限制	<ul> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as rename tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>Binlog日志:</li> <li>The binary logging feature must be enabled. The value of the binlog_format parameter must be set to row. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> <li>如为增量迁移任务, DTS要求源数据库的本地Binlog日志保存24小时以上,如为库表结构迁移和增量迁移任务, DTS要求源数据库的本地Binlog日志至少保留7天以上(您可在库表结构迁移完成 后将Binlog保存时间设置为24小时以上),否则DTS可能因无法获取Binlog而导致任务失败,极端情况下甚至可能会导致数据不一致或丢失。由于您所设置的Binlog日志保存时间低于DTS要求的时间进而导致的问题,不在DTS的SLA保障范围内。</li> <li>源库的操作限制: 在库表结构迁移阶段,请勿执行库或表结构变更的DDL操作,否则数据迁移任务会失败。</li> </ul>
其他限制	<ul> <li>不支持全量数据初始化,即DTS不会将源集群中迁移对象的存量数据迁移至目标DataHub实例。</li> <li>仅支持表级别的数据迁移。</li> <li>Read-only nodes of the source cluster cannot be migrated.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance after the task is resumed.</li> </ul>

## 2.3. Precautions and limits for migrating data from a PolarDB for Oracle cluster

This topic describes the precautions and limits when you migrate data from a PolarDB for Oracle cluster. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from a PolarDB for MySQL cluster

You can view the precautions and limits based on the following migration scenarios:

- Migrate data between PolarDB for Oracle clusters
- Migrate data from a PolarDB for Oracle cluster to a self-managed Oracle database

#### Migrate data between PolarDB for Oracle clusters

The following table describes the precautions and limits.

Category	Description		
	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as renaming tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrate the entire database.</li> </ul>		
	<ul> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>		
	• The write-ahead logging (WAL) feature must be enabled.		
Limits on the source database	<ul> <li>If you perform only incremental data migration, the WAL logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the WAL logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the WAL logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of WAL logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>		
	• Limits on operations:		
	<ul> <li>During schema migration and full data migration, do not perform data definition language (DDL) operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>		
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>		

Category	Description
	<ul> <li>A single data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each database.</li> <li>During incremental data migration, if you select a schema as the object to be migrated, take note of the following limits: If you create a table in the schema or run the RENAME command to rename the table, you must run the ALTER TABLE s chema.table REPLICA IDENTITY FULL; command before you write data to the table.</li> </ul>
	<b>Note</b> Replace the schema and table in the preceding sample command with the actual schema name and table name.
	• To ensure that the delay time of incremental data migration is accurate, DTS adds a heartbeat table named <a href="https://dts_postgres_heartbeat">dts_postgres_heartbeat</a> to the source database. The following figure shows the schema of the heartbeat table.
	execute(F8) Row Details Plan(F7) Format(F9)           Messages         Results1         Cross Database SQL Query           SLOT NAME         v         REVICE TIME         v         FLUSHED_LSN         v         UPDATE TIME         v         DTS_SERVICE TIME         v
Other limits	<ul> <li>1 voi 1005100402500 0/44 2020-03-25 10:55:47.655187:08 1055147.055187:08</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> </ul>
	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.
	• You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN , PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

#### Migrate data from a PolarDB for Oracle cluster to a self-managed Oracle database

The following table describes the precautions and limits.

Category

Category	Description	
	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as renaming tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrate the entire database.</li> </ul>	
	<ul> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>	
	<ul> <li>The write-ahead logging (WAL) feature must be enabled.</li> </ul>	
Limits on the source database	<ul> <li>If you perform only incremental data migration, the WAL logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the WAL logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the WAL logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of WAL logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>	
	Limits on operations:	
	<ul> <li>During full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>	
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Full Data Migration and Incremental Data Migration as the migration types.</li> </ul>	

Category	Description
	<ul> <li>Schema migration is not supported in this scenario. Before you configure a data migration task, you must create databases and tables in the destination instance.</li> <li>A single data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each database.</li> <li>During incremental data migration, if you select a schema as the object to be migrated, take note of the following limits: If you create a table in the schema or run the RENAME command to rename the table, you must run the ALTER TABLE s chema.table REPLICA IDENTITY FULL; command before you write data to the table.</li> </ul>
	<b>Note</b> Replace the schema and table in the preceding sample command with the actual schema name and table name.
Other limits	• To ensure that the delay time of incremental data migration is accurate, DTS adds a heartbeat table named <a href="https://dts_postgres_heartbeat">dts_postgres_heartbeat</a> to the source database. The following figure shows the schema of the heartbeat table.
	execute(F8)         Row Details         Plan(F7)         Format(F9)           1         select * from "dtstest"."dts_postgres_heartbeat"           Messages         Results1         Cross Database SQL Query           SLOT_NAME         *         REVICE_TIME         *         FLUSHED_LSN         *         UPDATE_TIME         *         DTS_SERVICE_TIME         *           1         v61         1585104942560         0/44         mm1         2020-03-25 10:55:47.556187408         198510494579
	• Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.
	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.
	<ul> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the PRECISION function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> </ul>
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

Category	Description	
Special cases	If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Connect an on-premises database to Alibaba Cloud and Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.	
	Notice When you configure the source Oracle database in the DTS console, you must enter the SCAN IP address of the Oracle RAC in the Database Endpoint or IP Address field.	

# 2.4. Limits for using PolarDB-X as the source instance

does not provide binary logs. To ensure data quality, we recommend that you take note of the limits when you perform business design, business development, and O&M changes.

#### Overview

- Limits related to business design
- Limits related to database architecture
- Limits related to O&M changes
- Potential risks related to data quality
- Suggestions to ensure data quality

#### Limits related to business design

- All tables must have primary keys. Otherwise, data inconsistency may occur (the destination database may contain duplicate data records).
- We recommend that you do not use the global secondary indexes (GSIs) of because they are updated asynchronously. If you use GSIs, DTS can guarantee only eventual consistency of data.
- The databases that you want to synchronize cannot be deployed in mixed mode (where the unit mode and the copy mode are mixed).

(?) Note In unit mode, users perform read and write operations in their respective unit nodes. Two-way synchronization is implemented between the databases in each unit node and those in the central node. In copy mode, users write data to the databases in the central node. The data is then synchronized to the databases in each unit node.

- If you use the underlying MySQL databases of to configure two-way data synchronization tasks, you must convert the Float and Double data types into Decimal for business tables. If you use as the source instance of a one-way data synchronization task, a data migration task, or a change tracking task, you do not need to convert the Float and Double data types.
- Data synchronization tasks, data migration tasks, and change tracking tasks do not support the following objects of : stored procedures, triggers, functions, views, and events.

- DTS does not support schema synchronization for . You must manually create objects such as databases and tables in the destination instance.
- The source instance must have sufficient capacity to support business growth.
- If MySQL databases of version 5.7 and 8.0 run on the instance, you cannot use the instance as the source instance of a change tracking task. You must configure a change tracking task for each MySQL database to track and consume data from the instance.

#### Limits related to database architecture

- The ApsaraDB RDS for MySQL instances used by a instance cannot be used by other instances.
- For data synchronization or migration between instances, the source and destination ApsaraDB RDS for MySQL instances must have equivalent deployment. For example, if the source instance uses four ApsaraDB RDS for MySQL instances, the destination instance must also use four ApsaraDB RDS for MySQL instances with the same specifications.
- The sharding rules of the source and destination instances must be the same. Otherwise, the data synchronization or migration task cannot be created.
- You can synchronize, migrate, or track the data of business tables in an instance. You cannot synchronize, migrate, or track the data of metadata tables or system tables in the instance.

Category	Description	Impact and solution
	Change a sharding rule, for example, change the shard key of a database or table or change the number of shards.	<ul> <li>Not supported. You must perform the following steps to recreate the task:</li> <li>1. Stop and delete the original DTS task.</li> <li>2. Wait until the changes are completed in the source database. Then, clear the data that has</li> </ul>
	Change the number of instances at the storage layer, for example, scale out instances and migrate frequently-accessed tables.	<ul><li>been synchronized or migrated to the destination database.</li><li>3. Configure a data synchronization or migration task for each ApsaraDB RDS for MySQL instance in the instance.</li></ul>
Storage layer	Change the specifications of instances and switch workloads at the storage layer.	DTS tasks are not affected.
	Change parameter settings.	The parameter settings of the source and destination databases must be the same. To change the parameter settings of instances at the storage layer, you must make sure that new parameters do not affect previous parameters.
		<b>Note</b> If you are not sure about the impact of changing parameter settings, you can contact technical support of Database Expert Service.

#### Limits related to O&M changes

Category	Description	Impact and solution
	Change backup and recovery policies, and enable auditing and diagnostics for instances at the storage layer.	The change takes effect only on the current instance and does not affect other instances with replication relationships.
DTS task	Perform DDL operations.	If you configure a DTS task to replicate data from multiple ApsaraDB RDS for MySQL instances in a instance to the destination database, performing DDL operations may cause task latency.
DDL operations at the database or table level	Add tables.	<ul> <li>Not supported. You must perform the following steps:</li> <li>1. After you perform DDL operations to create tables in the destination database, perform the same DDL operations to create tables in the source database.</li> <li>2. Add the new tables to the objects of the data synchronization task.</li> <li>You can write data to the source database only after the preceding operations are completed. If you select tables as the objects of the data synchronization task, you must add the new tables to the source and destination instances at the storage layer.</li> </ul>
	Add fields, add secondary indexes, delete indexes, and modify indexes (except for replacing secondary indexes with unique indexes).	<ul> <li>If you configure a DTS task based on a instance, you must perform DDL operations in the destination database and then perform the same DDL operations in the source database.</li> <li>If you configure a DTS task based on ApsaraDB RDS for MySQL instances attached to a instance, DTS automatically updates the operations. The following operations are supported: Add fields, add secondary indexes, delete indexes, and modify indexes (except for replacing secondary indexes with unique indexes).</li> </ul>
	Perform other DDL operations.	Only the preceding DDL operations are supported.

Category	Description	Impact and solution
	Perform a switchover.	Before you perform a switchover, make sure that the DTS task is not delayed. Otherwise, data quality issues occur.
	Perform a failover that meets the requirements of recovery point objective (RPO).	
Switchover	<b>?</b> Note RPO represents the maximum amount of data that can be lost after a recovery from a failure. RPO is measured by time.	If a failure (such as network interruption, equipment failure, or Internet data center failure) occurs and
NoteSwitchover:After you useDTS tosynchronizeor migratedata from thesourcedatabase tothedestinationdatabase, youswitchworkloadsfrom thesourcedatabase tothedestinationdatabase, youswitchworkloadsfrom thesourcedatabase tothedestination	Warning Failover: If the source instance or the data center where the source instance resides fails, you can switch workloads to a backup system. A failover is a lossy operation.	the DTS task is delayed, you may need to perform a failover. In this case, if the difference between the time when the last data entry is updated to the destination database and the time when the failure occurs is less than the RPO, you can perform a failover to recover your business. For example, if the RPO is 5 minutes, the quality of the data within the 5 minutes cannot be guaranteed after you perform a failover. You may need to revise the data to ensure consistency.
dat abase.	Perform a failover that does not meet the requirements of RPO.	The DTS task may be delayed because of the following reasons: a large number of DDL operations is performed in the source database, a network failure occurs, and the performance of the destination database is unfavorable. In this case, if the data center fails and the difference between the time when the last data entry is updated to the destination database and the time when the failure occurs is greater than the RPO, we recommend that you wait until the data center recovers before you perform a failover. For example, if the RPO is 5 minutes, the quality of the data within the 5 minutes cannot be guaranteed after you perform a failover. You may need to revise the data to ensure consistency.

#### Potential risks related to data quality

Some changes or switchover operations may cause data quality issues such as schema inconsistency between the source and destination databases.

- If data latency occurs between the primary and secondary databases of the source instance, the data written to the primary database is not updated to the secondary database in a timely manner. In this case, if you perform a primary/secondary switchover in the source instance, DTS uses the secondary database of the source instance as the source database for data synchronization, data migration, or change tracking. As a result, the data that is not updated to the secondary database is lost.
- If the DTS task is resumed from a network failure after you perform a switchover, DTS attempts to synchronize, migrate, or track the data generated before the failure occurs. This mechanism prevents data loss in the destination database. In this case, if the destination tables do not have primary keys, data will be inconsistent between the source and destination databases. If the destination tables have primary keys, data may not be consistent when DTS implements the retry mechanism, but data will remain consistent after the retry ends.
- The DTS task may be delayed due to network failures and DDL operations.
- The DTS task may be delayed or interrupted due to changes to the source database, unfavorable performance of the destination database, and schema inconsistency.

Alibaba Cloud cannot solve the preceding issues. You must recreate a DTS task or adjust the source and destination databases.

#### Suggestions to ensure data quality

- You must perform all DDL operations with caution. All DDL operations must be confirmed by the technical engineers to comply with the preceding limits.
- Do not directly perform DDL operations in your program code.

# 2.5. Solutions for using DRDS as the source instance

Data Transmission Service (DTS) allows you to synchronize or migrate data from a instance to the destination database. However, if you directly use a instance as the source instance of a DTS task and more than two ApsaraDB RDS for MySQL instances are attached to the DRDS instance, performance bottlenecks and stability risks may occur. Your business may be affected. To achieve higher performance and stability, we recommend that you configure a task for each ApsaraDB RDS for MySQL instance.

#### Prerequisites

You have read the Limits for using PolarDB-X as the source instance and have made sure to follow the relevant conventions and precautions.

#### Solutions

Solutio	Configuration method	Description
n	comiguation method	Description

Solutio n	Configuration method	Description
Solutio n 1	Configure a DTS task for each ApsaraDB RDS for MySQL instance that is attached to the instance. When you configure DTS tasks, you must map the database and table names of all the ApsaraDB RDS for MySQL instances to the database and table names of the destination instance. Compared with directly using the instance as the source instance of a DTS task, this solution ensures higher performance and stability.	We recommend that you use Solution 1, which provides higher performance and stability than Solution 2. In addition, the number of ApsaraDB RDS for MySQL instances attached to the instance is unlimited.
Solutio n 2	Directly use the instance as the source instance to configure a data synchronization or migration task. If more than two ApsaraDB RDS for MySQL instances are attached to the DRDS instance, the stability and operability of the DTS task may be affected. Your business may also be affected.	If you use Solution 2 and more than two ApsaraDB RDS for MySQL instances are attached to the source instance, the performance and stability of the DTS task may be compromised.

#### Comparison between Solution 1 and Solution 2

ltem	Solution 1	Solution 2
Performance	Multiple DTS tasks provide higher performance and support large amounts of data writes to the instance.	Only one DTS task is configured to migrate or synchronize data from the instance. When the business system writes large amounts of data to the source instance, performance bottlenecks occur.
Stability	High. Multiple DTS tasks are configured to migrate or synchronize data from the ApsaraDB RDS for MySQL instances attached to the Setting instance. If one of the DTS tasks fails, the other DTS tasks are not affected. You only need to recover the failed DTS task.	Medium. Only one DTS task is configured to migrate or synchronize data from the instance. If the DTS task fails, you must troubleshoot the entire task to resume data transmission.
Ease of use	You must configure multiple DTS tasks. You must configure database and table name mapping for each task. You must map the database and table names of multiple ApsaraDB RDS for MySQL instances in the source instance to the names of the databases and tables in the destination instance.	You need to configure only one DTS task for the source instance.
Resource usage	Multiple DTS instances are required.	Only one DTS instance is required.

# How do the two solutions implement synchronization or incremental migration of DDL operations

DTS does not support synchronization or incremental migration of DDL operations from a instance. If DDL operations are performed in the source instance during data synchronization or migration, you can perform the following steps to ensure that data can be written to the destination database:

- 1. Release the DTS task.
- 2. Clear the destination database.
- 3. Reconfigure the task.

In some scenarios where is the source instance, you can synchronize or migrate DDL operations without releasing the DTS task. The following table describes specific scenarios and corresponding operations.

Scenarios	Operation	
You use Solution 1 and select tables as the objects to synchronize.	<ul> <li>You can add tables to the objects that you select for the data synchronization task.</li> <li>You can add or remove columns only when you synchronize data between two instances. To do this, perform the following steps: <ol> <li>Add columns to or remove columns from the objects that you select for the data synchronization task.</li> <li>Add or remove columns in the destination database, and then perform the same operations in the source database. When DTS detects that a column already exists in the destination database, DTS ignores the error and does not display a write failure.</li> </ol> </li> </ul>	
You use Solution 2 and select an entire database as the object to be migrated or synchronized.	<ul> <li>You can only add tables. To do this, you must add tables in the destination database and then perform the same operations in the source database.</li> <li>You cannot add or remove columns.</li> <li>Warning If you add a column in the source instance, some physical tables at the underlying layer may contain the column whereas some physical tables do not contain the column. When DTS assembles the SQL statement, DTS may fail to find the column or lose the data in the column.</li> </ul>	
You use Solution 2 and do not select an entire database as the object to be migrated or synchronized.	<ul> <li>You can only add tables. To do this, perform the following steps: <ol> <li>Add tables to the objects that you select for the data synchronization task.</li> <li>Add tables in the destination database, and then perform the same operations in the source database.</li> </ol> </li> <li>You cannot add or remove columns.</li> </ul>	

# 2.6. Precautions and limits for migrating data from an Oracle database

This topic describes the precautions and limits when you migrate data from an Oracle database. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from an Oracle database

- Take note of precautions and limits in the following data migration scenarios:
- Migrate data from a self-managed Oracle database to a PolarDB for Oracle cluster
- Migrate data from a self-managed Oracle database to a MySQL database
- Migrate data from a self-managed Oracle database to a PolarDB for MySQL cluster
- Migrate data from a self-managed Oracle database to an AnalyticDB for PostgreSQL instance
- Migrate data from a self-managed Oracle database to a Message Queue for Apache Kafka instance or a self-managed Kafka cluster
- Migrate data between self-managed Oracle databases

## Migrate data from a self-managed Oracle database to a PolarDB for Oracle cluster

The following table describes the precautions and limits.

Category

Category	Description
	• The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.
	• If the source database is connected over Express Connect, you must specify a virtual IP address (VIP) for the database when you configure the source database information.
	• If the source database is an Oracle RAC database hosted on Elastic Compute Service (ECS) or connected over Express Connect, VPN Gateway, Smart Access Gateway, Database Gateway, or Cloud Enterprise Network (CEN), you can use a single VIP rather than a Single Client Access Name (SCAN) IP address when you configure the source database information. After you specify the VIP, node failover is not supported for the Oracle RAC database.
	Requirements for the objects to migrate:
	• The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
	<ul> <li>If the version number of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.</li> </ul>
	<ul> <li>If you select tables as the objects to migrate and you want to edit the tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
Limits on	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
the source	• The redo logging and archive logging must be enabled.
database	<ul> <li>For an incremental data migration task, redo logs and archive logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, redo logs and archive logs of the source database must be stored for at least seven days. After the full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, Data Transmission Service (DTS) may fail to obtain the redo logs and archive logs, and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of redo logs and archive logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>
	<ul> <li>During data migration, do not update LONGTEXT fields. Otherwise, the data migration task fails.</li> </ul>

Category	Description
Other limits	• Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.
	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads over to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.

# Migrate data from a self-managed Oracle database to a MySQL database

The following table describes the precautions and limits when you migrate data to MySQL databases, such as self-managed MySQL databases and ApsaraDB RDS for MySQL instances.

Category Description

Category	Description
	• The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.
	• If the source database is connected over Express Connect, you must specify a VIP for the database when you configure the source database information.
	• If the source database is an Oracle RAC database hosted on Elastic Compute Service (ECS) or connected over Express Connect, VPN Gateway, Smart Access Gateway, Database Gateway, or Cloud Enterprise Network (CEN), you can use a single VIP rather than a Single Client Access Name (SCAN) IP address when you configure the source database information. After you specify the VIP, node failover is not supported for the Oracle RAC database.
	Requirements for the objects to migrate:
	<ul> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> </ul>
	<ul> <li>If the version number of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.</li> </ul>
	<ul> <li>If you select tables as the objects to migrate and you want to edit the tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
l imits on	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
the source	• The redo logging and archive logging must be enabled.
database	<ul> <li>For an incremental data migration task, redo logs and archive logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, redo logs and archive logs of the source database must be stored for at least seven days. After the full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, Data Transmission Service (DTS) may fail to obtain the redo logs and archive logs, and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of redo logs and archive logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>
	<ul> <li>During data migration, do not update LONGTEXT fields. Otherwise, the data migration task fails.</li> </ul>

Category	Description
Ot her limit s	<ul> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads over to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.</li> </ul>
Special cases	<ul> <li>If the destination database runs on an ApsaraDB RDS for MySQL instance, take note of the following limits:</li> <li>Table names in the ApsaraDB RDS for MySQL instance are case-insensitive. If a table name in the source Oracle database contains uppercase letters, ApsaraDB RDS for MySQL converts all uppercase letters to lowercase letters before a table is created.</li> <li>If the source Oracle database contains identical table names that differ only in capitalization, these table names are identified as duplicate. As a result, the "The object already exists" message may be displayed during schema migration. To prevent name conflicts in the destination database, you can use the object name mapping feature to capitalize the table names. For more information, see Object name mapping.</li> <li>DTS automatically creates a destination database is invalid, you must create a database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must create a database in the ApsaraDB RDS for MySQL instance.</li> </ul>

# Migrate data from a self-managed Oracle database to a PolarDB for MySQL cluster

The following table describes the precautions and limits when you migrate data from a self-managed Oracle database to a PolarDB for MySQL cluster.

Category	Description
----------	-------------

Category	Description
	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>If the source database is connected over Express Connect you must specify a VID for the source database is connected over Express Connect you must specify a VID for the source database belongs and the source database is connected over Express Connect you must specify a VID for the source database belongs and the source database belongs are specified over Express Connected over Express Connected</li></ul>
	database when you configure the source database information.
	• If the source database is an Oracle RAC database hosted on Elastic Compute Service (ECS) or connected over Express Connect, VPN Gateway, Smart Access Gateway, Database Gateway, or Cloud Enterprise Network (CEN), you can use a single VIP rather than a Single Client Access Name (SCAN) IP address when you configure the source database information. After you specify the VIP, node failover is not supported for the Oracle RAC database.
	Requirements for the objects to migrate:
	<ul> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> </ul>
	<ul> <li>If the version number of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.</li> </ul>
	<ul> <li>If you select tables as the objects to migrate and you want to edit the tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
Limits on the source database	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
	• The redo logging and archive logging must be enabled.
	<ul> <li>For an incremental data migration task, redo logs and archive logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, redo logs and archive logs of the source database must be stored for at least seven days. After the full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, Data Transmission Service (DTS) may fail to obtain the redo logs and archive logs, and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of redo logs and archive logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema</li> </ul>
	<ul> <li>migration, full data migration, and incremental data migration as the migration types.</li> <li>During data migration, do not update LONGTEXT fields. Otherwise, the data migration task fails.</li> </ul>

Category	Description
Ot her limit s	<ul> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads over to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.</li> </ul>
Special cases	<ul> <li>If the destination database runs on a PolarDB for MySQL cluster, take note of the following limits:</li> <li>Table names in the PolarDB for MySQL cluster are case-insensitive. If a table name in the source Oracle database contains uppercase letters, PolarDB for MySQL converts all uppercase letters to lowercase letters before a table is created.</li> <li>If the source Oracle database contains identical table names that differ only in capitalization, these table names are identified as duplicate. During schema migration, the following message is returned: "The object already exists". To prevent name conflicts in the destination database, you can use the object name mapping feature to capitalize the table names. For more information, see Object name mapping.</li> <li>DTS automatically creates a destination database in the PolarDB for MySQL cluster. However, if the name of the source database is invalid, you must create a database in the PolarDB for MySQL cluster before you configure the data migration task. For more information, see Database Management.</li> </ul>

# Migrate data from a self-managed Oracle database to an AnalyticDB for PostgreSQL instance

The following table describes the precautions and limits when you migrate data from a self-managed Oracle database to an AnalyticDB for PostgreSQL instance.

Category Description

Category	Description	
Limits on the source database	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>If the source database is connected over Express Connect, you must specify a VIP for the database when you configure the source database information.</li> </ul>	
	<ul> <li>If the source database is an Oracle RAC database hosted on Elastic Compute Service (ECS) or connected over Express Connect, VPN Gateway, Smart Access Gateway, Database Gateway, or Cloud Enterprise Network (CEN), you can use a single VIP rather than a Single Client Access Name (SCAN) IP address when you configure the source database information. After you specify the VIP, node failover is not supported for the Oracle RAC database.</li> </ul>	
	Requirements for the objects to migrate:	
	<ul> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> </ul>	
	• If the version number of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.	
	<ul> <li>If you select tables as the objects to migrate and you want to edit the tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>	
	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>	
	• The redo logging and archive logging must be enabled.	
	<ul> <li>For an incremental data migration task, redo logs and archive logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, redo logs and archive logs of the source database must be stored for at least seven days. After the full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, Data Transmission Service (DTS) may fail to obtain the redo logs and archive logs, and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of redo logs and archive logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>	
	• Limits on operations:	
	• During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.	
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>	
	<ul> <li>During data migration, do not update LONGTEXT fields. Otherwise, the data migration task fails.</li> </ul>	
Category	Description	
-------------------	--	--
Ot her limit s	• Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.	
	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.	
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads over to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.	

#### Migrate data from a self-managed Oracle database to a Message Queue for Apache Kafka instance or a self-managed Kafka cluster

```
Category
```

Category	Description
	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>If the source database is connected over Express Connect, you must specify a VIP for the database when you configure the source database information.</li> </ul>
	• If the source database is an Oracle RAC database hosted on Elastic Compute Service (ECS) or connected over Express Connect, VPN Gateway, Smart Access Gateway, Database Gateway, or Cloud Enterprise Network (CEN), you can use a single VIP rather than a Single Client Access Name (SCAN) IP address when you configure the source database information. After you specify the VIP, node failover is not supported for the Oracle RAC database.
	Requirements for the objects to migrate:
	<ul> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> </ul>
	<ul> <li>If the version number of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.</li> </ul>
Limits on	<ul> <li>If you select tables as the objects to migrate and you want to edit the tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
the source database	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
	• The redo logging and archive logging must be enabled.
	<ul> <li>For an incremental data migration task, redo logs and archive logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, redo logs and archive logs of the source database must be stored for at least seven days. After the full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, Data Transmission Service (DTS) may fail to obtain the redo logs and archive logs, and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of redo logs and archive logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>
	<ul> <li>During data migration, do not update LONGTEXT fields. Otherwise, the data migration task fails.</li> </ul>

Category	Description
Ot her limit s	<ul> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.</li> </ul>
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads over to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.

#### Migrate data between self-managed Oracle databases

The following table describes the precautions and limits when you migrate data between selfmanaged Oracle databases.

Category	Description
	• The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.
	• If the source database is connected over Express Connect, you must specify a VIP for the database when you configure the source database information.
	• If the source database is an Oracle RAC database hosted on Elastic Compute Service (ECS) or connected over Express Connect, VPN Gateway, Smart Access Gateway, Database Gateway, or Cloud Enterprise Network (CEN), you can use a single VIP rather than a Single Client Access Name (SCAN) IP address when you configure the source database information. After you specify the VIP, node failover is not supported for the Oracle RAC database.
	Requirements for the objects to migrate:
	<ul> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> </ul>
	<ul> <li>If the version number of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.</li> </ul>
	<ul> <li>If you select tables as the objects to migrate and you want to edit the tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	<ul> <li>If you want to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
Limits on	• The redo logging and archive logging must be enabled.
the source database	<ul> <li>For an incremental data migration task, redo logs and archive logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration task, redo logs and archive logs of the source database must be stored for at least seven days. After the full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, Data Transmission Service (DTS) may fail to obtain the redo logs and archive logs, and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of redo logs and archive logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>
	Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.</li> </ul>
	<ul> <li>During data migration, do not update LONGTEXT fields. Otherwise, the data migration task fails.</li> </ul>

6.1	Provide the second s	
Category	Description	
Ot her limit s	<ul> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> </ul>	
	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.	
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads over to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.	

### 2.7. Precautions and limits for migrating data from a PostgreSQL database

This topic describes the precautions and limits when you migrate data from a PostgreSQL database, such as a self-managed PostgreSQL database and an ApsaraDB RDS for PostgreSQL instance. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from a PostgreSQL database

Take note of precautions and limits in the following data migration scenarios:

- Migrate data between PostgreSQL databases
- Migrate data from a PostgreSQL database to a MySQL database
- Migrate data from a self-managed PostgreSQL database to a PolarDB for Oracle cluster

#### Migrate data between PostgreSQL databases

• Migrate data between ApsaraDB RDS for PostgreSQL instances

Category

Category	Description
	<ul> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records. The name of the source database cannot contain hyphens (-). Example: dts-testdata.</li> <li>If you select tables as objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>If you need to migrate incremental data, you must make sure that the following</li> </ul>
	requirements are met:
	The value of the wal_level parameter must be set to <i>logical</i> .
Limits on the source database	For an incremental data migration, the WAL logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration, the WAL logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the WAL logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of WAL logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability or performance.
	• Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select schema migration, full data migration, and incremental data migration as the migration types.
	<ul> <li>If you need to perform a primary/secondary switchover on the source ApsaraDB RDS for PostgreSQL instance, the Logical Replication Slot Failover feature must be enabled. This prevents logical subscriptions from being interrupted and ensures that your data migration task can run as expected. For more information, see Logical Replication Slot Failover.</li> <li>A data migration task can migrate data from only a single database. To migrate data from multiple databases, you must create a data migration task for each database.</li> </ul>

Category	<ul> <li>During incremental data migration, if you select a schema as the object to Description migrate, take note of the following limits: If you create a table in the schema or execute the RENAME statement to rename the table, you must execute the ALTE</li> </ul>
	R TABLE schema.table REPLICA IDENTITY FULL; statement before you write data to the table.
	<b>Note</b> Replace the schema and table in the preceding sample statement with your actual schema name and table name.
	• DTS does not check the validity of metadata such as sequences. You must manually check the validity of metadata.
	<ul> <li>After your workloads are switched to the destination database, newly written sequences do not increment from the maximum value of the sequences in the source database. Therefore, you must query the maximum value of the sequences in the source database before you switch your workloads to the destination database. Then, you must specify the queried maximum value as the starting value of the sequences in the destination database. You can execute the following statements to query the maximum value of the sequences in the source database:</li> </ul>
	do language plpgsql \$\$
	declare nsp name;
	rel name;
	val int8;
	begin
	<pre>for nsp,rel in select nspname,relname from pg_class t2 , pg_namespace t3 where t2.relnamespace=t3.oid and t2.relkind='S'</pre>
	<pre>ioop     execute format(\$_\$select last_value from %I.%I\$_\$, nsp, rel) into val:</pre>
	raise notice '%',
	<pre>format(\$_\$select setval('%1.%1'::regclass, %s);\$_\$, nsp, rel, val+1);</pre>
	end loop;
	end; \$\$;
	• DTS creates the following temporary tables in the source database to obtain the DDL statements of incremental data, the schemas of incremental tables, and the heartbeat information. During data migration, do not delete temporary tables in the source database. Otherwise, the data migration task may fail. After the DTS instance is released, temporary tables are automatically deleted.
	<pre>public.DTS_PG_CLASS , public.DTS_PG_ATTRIBUTE , public.DTS_PG_TY PE , public.DTS_PG_ENUM , public.DTS_POSTGRES_HEARTBEAT , public DTS_DDL_COMMAND .and public.DTS_ARGS_SESSION .</pre>
	<ul> <li>To ensure that the latency of incremental data migration is accurate, DTS creates a heartbeat table named dts postgres heartbeat in the source database.</li> </ul>
Other limits	

Category	<ul> <li>During incremental data migration, DTS creates a replication slot for the source Description database. The replication slot is prefixed with dts_sync_</li> <li>DTS automatically</li> </ul>
	clears historical replication slots every 90 minutes to reduce storage usage.
	<ul> <li>Note</li> <li>After the DTS instance is released, the replication slot is automatically deleted. If you modify the password of the source database or delete the IP address whitelist of DTS, the replication slot cannot be automatically deleted. In that case, you must delete it from the source database to prevent it from piling up.</li> <li>If the data migration task is released or fails, DTS automatically clears the replication slot. If a primary/secondary switchover is performed on the source ApsaraDB RDS for PostgreSQL instance, you must log on to the secondary instance to clear the replication slot.</li> </ul>
	Query Editor Query History Scratch Pad 🗙
	1 SELECT * FROM pg_replication_slots;
	Data Output Explain Messages Notifications
	slot_name plugin name slot_type text oid atabase of text oid atabase ative boolean ative
	1 dts_sync_ohu pgoutput logical 16 dtstestdata false true
	<ul> <li>performance of the source and destination databases. We recommed that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of tablespace used by the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLU MN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the data in the destination instance after the task is resumed.</li> </ul>

Description
During data migration, do not modify the endpoint and zone of the source ApsaraDB RDS for PostgreSQL instance. Otherwise, the data migration task fails.

• Migrate data from a self-managed PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance

ategory
---------

Category	Description
Limits on the source database	<ul> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>The name of the source database cannot contain hyphens (-). Example: dtstestdata.</li> <li>If you select tables as objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run at task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> <li>The value of the wal_level parameter must be set to <i>logical</i>.</li> <li>For an incremental data migration, the WAL logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. For a full data and incremental data migration, the WAL logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the WAL logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of WAL logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability or performance.</li> <li>Limits on operations:</li> <li>If you perform a migration and full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> <li>Dur</li></ul>
	<ul> <li>Data may be inconsistent between the primary and secondary nodes of the source database due to migration latency. Therefore, you must use the primary node as the data source when you migrate data.</li> <li>A data migration task can migrate data from only a single database. To migrate data from multiple databases, you must create a data migration task for each database.</li> </ul>

Category	<ul> <li>During incremental data migration, if you select a schema as the object to Description migrate, take note of the following limits: If you create a table in the schema or</li> </ul>
	execute the RENAME statement to rename the table, you must execute the ALTE R TABLE schema.table REPLICA IDENTITY FULL; statement before you write data to the table.
	<b>Note</b> Replace the schema and table in the preceding sample statement with your actual schema name and table name.
	<ul> <li>DTS does not check the validity of metadata such as sequences. You must manually check the validity of metadata.</li> <li>After your workloads are switched to the destination database, newly written sequences do not increment from the maximum value of the sequences in the source database. Therefore, you must query the maximum value of the sequences in the source database before you switch your workloads to the destination database. Then, you must specify the queried maximum value as the starting value of the sequences in the destination database. You can execute the following statements to query the maximum value of the sequences in the source database:</li> </ul>
	<pre>do language plpgsql \$\$ declare   nsp name;   rel name;   val int8; begin   for nsp,rel in select nspname,relname from pg_class t2 ,   pg_namespace t3 where t2.relnamespace=t3.oid and t2.relkind='S'   loop     execute format(\$_\$select last_value from %I.%I\$_\$, nsp, rel) into val;     raise notice '%',     format(\$_\$select setval('%I.%I'::regclass, %s);\$_\$, nsp, rel, val+1);     end loop; end;     \$\$;</pre>
Other limits	<ul> <li>DTS creates the following temporary tables in the source database to obtain the DDL statements of incremental data, the schemas of incremental tables, and the heartbeat information. During data migration, do not delete temporary tables in the source database. Otherwise, the data migration task may fail. After the DTS instance is released, temporary tables are automatically deleted.</li> <li>public.DTS_PG_CLASS , public.DTS_PG_ATTRIBUTE , public.DTS_PG_TY</li> <li>PE , public.DTS PG ENUM , public.DTS POSTGRES HEARTBEAT , public</li> </ul>
	<ul> <li>.DTS_DDL_COMMAND , and public.DTS_ARGS_SESSION .</li> <li>To ensure that the latency of incremental data migration is accurate, DTS creates a heartbeat table named dts_postgres_heartbeat in the source database.</li> </ul>

Category	<ul> <li>During incremental data migration, DTS creates a replication slot for the source Description database. The replication slot is prefixed with <a href="https://dts_sync_">dts_sync_</a>. DTS automatically</li> </ul>
	clears historical replication slots every 90 minutes to reduce storage usage.
	<b>Note</b> If the data migration task is released or fails, DTS automatically clears the replication slot. If a primary/secondary switchover is performed on the source ApsaraDB RDS for PostgreSQL instance, you must log on to the secondary instance to clear the replication slot.
	Query Editor Query History Scratch Pad 🗙
	1 SELECT * FROM pg_replication_slots;
	Data Output Explain Messages Notifications
	alot_name plugin name biot_type datoid database have boolean active boolean activ
	1 dts_sync_ohu pgoutput logical 16 dtstestdata false true
	<ul> <li>performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of tablespace used by the destination database is larger than that of the source database.</li> </ul>
	<ul> <li>You must make sure that the precision settings for columns of the FLOAT or</li> <li>DOUBLE data type meets your business requirements. DTS uses the POUND (COLUMN)</li> </ul>
	MN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the
	FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.
	<ul> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.</li> </ul>

#### Migrate data from a PostgreSQL database to a MySQL database

The new DTS console supports the following migration scenarios:

- Migrate data from an ApsaraDB RDS for PostgreSQL instance to an ApsaraDB RDS for MySQL instance
- Migrate data from a self-managed PostgreSQL database to a self-managed MySQL database

The following table describes the precautions and limits.

Category	Description
	• The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.
	• The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
	The name of the source database cannot contain hypnens (-). Example: dts- testdata.
	<ul> <li>If you select tables as objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	• If you need to migrate incremental data, you must make sure that the following requirements are met:
	• The value of the wal_level parameter must be set to <i>logical</i> .
Limits on the source database	<ul> <li>For an incremental data migration, the WAL logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration, the WAL logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the WAL logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of WAL logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability or performance.</li> </ul>
	Limits on operations:
	<ul> <li>If you perform a primary/secondary switchover on the self-managed PostgreSQL database, the data migration task fails.</li> </ul>
	<ul> <li>If you need to perform a primary/secondary switchover on the source ApsaraDB RDS for PostgreSQL instance, the Logical Replication Slot Failover feature must be enabled. This prevents logical subscriptions from being interrupted and ensures that your data migration task can run as expected. For more information, see Logical Replication Slot Failover.</li> </ul>
	<ul> <li>During full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select full data migration and incremental data migration as the migration types.</li> </ul>
	• Before you configure a data migration task, you must create databases and tables in the destination instance.
	• A data migration task can migrate data from only a single database. To migrate data from multiple databases, you must create a data migration task for each database.
	• During incremental data migration, if you select a schema as the object to migrate, take note of the following limits: If you create a table in the schema or execute the

Category	RENAME statement to rename the table, you must execute the ALTER TABLE sch Description ema.table REPLICA IDENTITY FULL; statement before you write data to the
	table.
	<b>Note</b> Replace the schema and table in the preceding sample statement with your actual schema name and table name.
	<ul> <li>DTS creates the following temporary tables in the source database to obtain the DDL statements of incremental data, the schemas of incremental tables, and the heartbeat information. During data migration, do not delete temporary tables in the source database. Otherwise, the data migration task may fail. After the DTS instance is released, temporary tables are automatically deleted.</li> <li>public.DTS_PG_CLASS , public.DTS_PG_ATTRIBUTE , public.DTS_PG_TYP</li> <li>E , public.DTS_PG_ENUM , public.DTS_POSTGRES_HEARTBEAT , public.DT</li> <li>S_DDL_COMMAND , and public.DTS_ARGS_SESSION .</li> </ul>
	• To ensure that the latency of incremental data migration is accurate, DTS creates a heartbeat table named dts_postgres_heartbeat in the source database.
	<ul> <li>During incremental data migration, DTS creates a replication slot for the source database. The replication slot is prefixed with <a href="https://dts_sync_">dts_sync_</a>. DTS automatically clears historical replication slots every 90 minutes to reduce storage usage.</li> </ul>
Other limits	<ul> <li>Note</li> <li>After the DTS instance is released, the replication slot is automatically deleted. If you modify the password of the source database or delete the IP address whitelist of DTS, the replication slot cannot be automatically deleted. In that case, you must delete it from the source database to prevent it from piling up.</li> <li>If the data migration task is released or fails, DTS automatically clears the replication slot. If a primary/secondary switchover is performed on the source ApsaraDB RDS for PostgreSQL instance, you must log on to the secondary instance to clear the replication slot.</li> </ul>
	<ul> <li>database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN , PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT</li> </ul>

Category	<ul> <li>data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>Description</li> <li>DTS attempts to resume data migration tasks that failed within the last seven</li> </ul>
	days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.
Special cases	<ul> <li>During data migration, do not modify the endpoint and zone of the source ApsaraDB RDS for PostgreSQL instance. Otherwise, the data migration task fails.</li> <li>If you migrate data to an ApsaraDB RDS for MySQL instance, DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task. For more information, see Create a database on an ApsaraDB RDS for MySQL instance.</li> </ul>

## Migrate data from a self-managed PostgreSQL database to a PolarDB for Oracle cluster

The following table describes the precautions and limits.

Category

Category	Description
Category Limits on the source database	<ul> <li>Description</li> <li>The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.</li> <li>The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records. <ul> <li>The name of the source database cannot contain hyphens (-). Example: dts-testdata.</li> </ul> </li> <li>If you select tables as objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>If you need to migrate incremental data, you must make sure that the following requirements are met: <ul> <li>The value of the wal_level parameter must be set to <i>logical</i>.</li> <li>For an incremental data migration, the WAL logs of the source database must be stored for more than 24 hours. For a full data and incremental data migration, the WAL logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. OTS may fail to obtain the WAL logs and the task may</li> </ul></li></ul>
source database	the WAL logs of the source database must be stored for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the WAL logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of WAL logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability or performance.
	<ul> <li>Limits on operations:</li> <li>If you perform a primary/secondary switchover on the self-managed PostgreSQL</li> </ul>
	database, the data migration task fails.
	<ul> <li>During full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases occurs. To ensure data consistency, we recommend that you select full data migration and incremental data migration as the migration types.</li> </ul>
	<ul> <li>Before you configure a data migration task, you must create databases and tables in the destination instance.</li> </ul>
	• During incremental data migration, if you select a schema as the object to migrate, take note of the following limits: If you create a table in the schema or execute the RENAME statement to rename the table, you must execute the ALTER TABLE schema.table REPLICA IDENTITY FULL; statement before you write data to the table.
	<b>Note</b> Replace the schema and table in the preceding sample statement with your actual schema name and table name.
	• DTS creates the following temporary tables in the source database to obtain the

Category	DDL statements of incremental data, the schemas of incremental tables, and the Description heartbeat information. During data migration, do not delete temporary tables in
	the source database. Otherwise, the data migration task may fail. After the DTS instance is released, temporary tables are automatically deleted.
	public.DTS_PG_CLASS , public.DTS_PG_ATTRIBUTE , public.DTS_PG_TYP
	E , public.DTS_PG_ENUM , public.DTS_POSTGRES_HEARTBEAT , public.DT S_DDL_COMMAND , and public.DTS_ARGS_SESSION .
	• To ensure that the latency of incremental data migration is accurate, DTS creates a heartbeat table named dts postgres heartbeat in the source database.
	• During incremental data migration, DTS creates a replication slot for the source database. The replication slot is prefixed with <a href="https://dts_sync_">dts_sync_</a> . DTS automatically clears historical replication slots every 90 minutes to reduce storage usage.
	<b>Note</b> If the data migration task is released or fails, DTS automatically clears the replication slot. If a primary/secondary switchover is performed on the source ApsaraDB RDS for PostgreSQL instance, you must log on to the secondary instance to clear the replication slot.
Other limits	Query Editor Query History Scratch Pad 🗙
	1 SELECT * FROM pg_replication_slots;
	Data Output Explain Messages Notifications
	slot_name name plugin ame slot_type datoid oid atabase temporary boolean active boolean active_bid integer
	1 dts_sync_ohu pgoutput logical 16 dtstestdata false true
	<ul> <li>A data migration task can migrate data from only a single database. To migrate data from multiple databases, you must create a data migration task for each database</li> </ul>
	<ul> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> </ul>
	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of tablespace used by the destination database is larger than that of the source database.
	<ul> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN , PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> </ul>
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also execute the REVOKE statement to revoke the write permissions from the accounts used by DTS to access the destination instance. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.

### 2.8. Precautions and limits for migrating data from an SQL Server database

This topic describes the precautions and limits when you migrate data from an SQL Server database, such as a self-managed SQL Server database and an database. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from an SQL Server database

You can view the precautions and limits based on the following migration scenarios:

- Migrate data between SQL Server databases
- Migrate data from an SQL Server database to an AnalyticDB for MySQL cluster
- Migrate data from an SQL Server database to an AnalyticDB for PostgreSQL instance

#### Migrate data between SQL Server databases

Category

Category	Description
Limits on the source database	• Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.
	• The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
	<ul> <li>If you select tables as the objects to be migrated and you need to edit tables (such as renaming tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	• A single data migration task can migrate up to 10 databases. If you need to migrate more than 10 databases, we recommend that you split the tables to be migrated or configure multiple tasks to migrate the databases. Otherwise, the performance and stability of your data migration task may be compromised.
	<ul> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
	<ul> <li>The data logging feature must be enabled. The backup mode must be set to Full, and full logical backup must be performed.</li> </ul>
	<ul> <li>If you perform only incremental data migration, the data logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the data logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the data logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of data logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform data definition language (DDL) operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>

Category	Description
Other limits	<ul> <li>DTS does not migrate data of the following types: CURSOR, ROWVERSION, SQL_VARIANT, HIERACHYID, and GEOMETRY.</li> <li>If you set SQL Server Incremental Synchronization Mode to Incremental Synchronization Based on Logs of Source Database in the Configure Objects and Advanced Settings step, the tables to be migrated must have clustered indexes that contain primary key columns.</li> <li>If you migrate data between different versions of databases, make sure that the database versions are compatible.</li> <li>To ensure that the delay time of incremental data migration is accurate, DTS adds a heartbeat table to the self-managed SQL Server database. The name of the heartbeat table is Source table name_dts_mysql_heartbeat .</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the data in the destination instance after the task is resumed.</li> </ul>
Special cases	If the destination instance is an ApsaraDB RDS for SQL Server instance, take note of the following limits: DTS automatically creates a destination database in the ApsaraDB RDS for SQL Server instance. However, if the name of the source database is invalid, you must create a database in the ApsaraDB RDS for SQL Server instance before you configure the data migration task. For more information, see Create a database on an ApsaraDB RDS for SQL Server instance.

## Migrate data from an SQL Server database to an AnalyticDB for MySQL cluster

Category

Category	Description
	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as renaming tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> </ul>
	<ul> <li>A single data migration task can migrate up to 10 databases. If you need to migrate more than 10 databases, we recommend that you split the tables to be migrated or configure multiple tasks to migrate the databases. Otherwise, the performance and stability of your data migration task may be compromised.</li> </ul>
	<ul> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>
Limits on	<ul> <li>The data logging feature must be enabled. The backup mode must be set to Full, and full logical backup must be performed.</li> </ul>
the source database	<ul> <li>If you perform only incremental data migration, the data logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the data logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the data logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of data logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>
	• Limits on operations:
	<ul> <li>During schema migration and full data migration, do not perform data definition language (DDL) operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
	<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>

Category	Description
Other limits	<ul> <li>DTS does not migrate data of the following types: TIMESTAMP, CURSOR, ROWVERSION, HIERACHYID, SQL_VARIANT, SPATIAL GEOMETRY, SPATIAL GEOGRAPHY, and TABLE.</li> <li>If you set SQL Server Incremental Synchronization Mode to Incremental Synchronization Based on Logs of Source Database in the Configure Objects and Advanced Settings step, the tables to be migrated must have clustered indexes that contain primary key columns.</li> <li>To ensure that the delay time of incremental data migration is accurate, DTS adds a heartbeat table to the self-managed SQL Server database. The name of the heartbeat table is Source table name_dts_mysql_heartbeat</li> <li>Due to the limits of , if the disk space usage of the nodes in an cluster reaches 80%, the performance of data writing to the destination database is compromised and the DTS task is delayed. If the usage reaches 90%, data cannot be written to the destination database and error messages are returned. We recommend that you estimate the required disk space based on the objects that you want to migrate. You must ensure that the destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, concurrent INSERT Operations cause fragmentation in the tables of the destination databases. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DUBLE data type meets your business requirements. DTS uses the <u>ROUND (COLUMN, PRECISION)</u> function to retrieve values from columns of the FLOAT or DUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DUBLE data type to 380 digits.</li> <li>DTS attempts to resume data migration tasks that faile</li></ul>

## Migrate data from an SQL Server database to an AnalyticDB for PostgreSQL instance

Category De

Category	Description
Limits on the source database	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as renaming tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>A single data migration task can migrate up to 10 databases. If you need to migrate more than 10 databases, we recommend that you split the tables to be migrated or configure multiple tasks to migrate the databases. Otherwise, the performance and stability of your data migration task may be compromised.</li> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> <li>The data logging feature must be enabled. The backup mode must be set to Full, and full logical backup must be performed.</li> <li>If you perform only incremental data migration, the data logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. If you perform both full data logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of data logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> </ul>
	<ul> <li>Limits on operations:</li> <li>During schema migration and full data migration, do not perform data definition language (DDL) operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>
Limits on the source database	<ul> <li>The data logging feature must be enabled. The backup mode must be set to Full, and full logical backup must be performed.</li> <li>If you perform only incremental data migration, the data logs of the source database must be stored for more than 24 hours. If you perform both full data migration and incremental data migration, the data logs of the source database must be stored for a least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the data logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of data logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> <li>Limits on operations:</li> <li>During schema migration and full data migration, do not perform data definition language (DDL) operations to change the schemas of databases or tables. Otherwise, the data migration. Otherwise, data will be inconsistent between the source database durin data migration. Otherwise, data will be inconsistent between the source and destinatic databases. To ensure data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration as the migration types.</li> </ul>

Category	Description
	<ul> <li>DTS does not migrate data of the following types: TIMESTAMP, CURSOR, ROWVERSION, HIERACHYID, SQL_VARIANT, SPATIAL GEOMETRY, SPATIAL GEOGRAPHY, and TABLE.</li> <li>If you set SQL Server Incremental Synchronization Mode to Incremental Synchronization Based on Logs of Source Database in the Configure Objects and Advanced Settings step, the tables to be migrated must have clustered indexes that contain primary key columns.</li> </ul>
	• To ensure that the delay time of incremental data migration is accurate, DTS adds a heartbeat table to the self-managed SQL Server database. The name of the heartbeat table is Source table name_dts_mysql_heartbeat .
	• Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.
Other limits	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.
	• You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

## 2.9. Precautions and limits for migrating data from a Redis database

This topic describes the precautions and limits that you must take note of when you migrate data from a Redis database, such as a self-managed Redis database and an ApsaraDB for Redis database. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from a Redis database

You can view the precautions and limits based on the following migration scenarios: Migrate data between Redis databases

#### Migrate data between Redis databases

Category	Description
Limits on the source database	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>If you perform only incremental data migration, you must enable the data logging feature. In addition, the append-only file (AOF) logs of the source database must be stored for more than 24 hours.</li> <li>If you perform both full data migration and incremental data migration, you must enable the data logging feature. In addition, the AOF logs of the source database must be stored for at least seven days. After full data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the AOF logs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of AOF logs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> <li>Limits on operations: If you perform only full data migration as the migration types.</li> <li>Limits on migrating a standalone ApsaraDB for Redis instance to an ApsaraDB for Redis cluster instance: Each command can be run only on a single slot in an ApsaraDB for Redis cluster instance. If you perform operations on multiple keys in the source database and the keys belong to different slots, an error occurs.</li> <li>CROSSSLOT Keys in request don't hash to the same slot</li> </ul>

Category	Description
	• ApsaraDB for Redis Enhanced Edition instances (storage-optimized instances) cannot be used as the source and destination databases.
	<ul> <li>If the data eviction policy (maxmemory-policy) of the destination database is not set to noeviction, data may become inconsistent between the source and destination databases.</li> </ul>
	• If you run the EVAL or EVALSHA command to call Lua scripts, DTS cannot identify whether these Lua scripts are executed on the destination database. During incremental data migration, the destination database does not explicitly return the execution results of Lua scripts.
	<ul> <li>When you run the PSYNC or SYNC command to transfer data of the LIST type, DTS does not perform the flush operation on the existing data. Therefore, the destination database may contain duplicate data records.</li> </ul>
	• If an expiration policy is enabled for some keys in the source database, these keys may not be deleted in a timely manner after they expired. Therefore, the number of keys in the destination database may be less than that in the source database. You can run the info command to view the number of keys in the destination database.
	<b>Note</b> The number of keys that do not have an expiration policy or have not expired is the same in the source and destination databases.
Other	• To ensure compatibility, the version of the destination Redis database must be the same as or later than the version of the source Redis database.
limits	<b>Note</b> If the version of the destination database is earlier than the version of the source database, database compatibility issues may occur.
	• Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.
	• During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the REVOKE command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.
	• If both the source and destination databases are ApsaraDB for Redis Community Edition or Enhanced Edition, take note of the following limits: If the database specifications are changed (for example, the specifications are upgraded or the port number is changed), DTS cannot obtain continuous log data and correct connection information. As a result, the data migration task is interrupted. To ensure data consistency, we recommend that you delete the data in the destination database and reconfigure the data migration task after the database specifications are changed.

Category	Description
Special cases	<ul> <li>If the source database is a self-managed Redis database, take note of the following limits:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no data manipulation language (DML) operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the migration task is too high, you can perform a DML operation on the source database to update the latency.</li> <li>Note If you select an entire database as the object to be migrated, you can create a heartbeat table. The heartbeat table is updated or receives data every second.</li> </ul>

## 2.10. Precautions and limits for migrating data from a MongoDB database

This topic describes the precautions and limits when you migrate data from a MongoDB database, such as a self-managed MongoDB database and an ApsaraDB for MongoDB database. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Scenarios of migrating data from a MongoDB database

You can view the precautions and limits based on the following migration scenarios:

- Migrate data from a MongoDB database (standalone architecture) to another MongoDB database (standalone, replica set, or sharded cluster architecture)
- Migrate data from a MongoDB database (replica set architecture) to another MongoDB database (replica set architecture or sharded cluster architecture)

#### Migrate data from a MongoDB database (standalone architecture) to another MongoDB database (standalone, replica set, or sharded cluster architecture)

Category	Description
Limits on the source database	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select collections as the objects to be migrated and you need to edit collections (such as renaming collections), up to 1,000 collections can be migrated in a single data migration task. If you run a task to migrate more than 1,000 collections, a request error occurs. In this case, we recommend that you split the collections to be migrated the entire database.</li> <li>Limits on operations:</li> <li>During schema migration and full data migration, do not change the schemas of databases or collections. Otherwise, the data migration task fails.</li> <li>Incremental data migration is not supported in this scenario. To ensure data consistency, we recommend that you do not write data to the source MongoDB database during full data migration.</li> </ul>

Category	Description
Other limits	<ul> <li>Only schema migration and full data migration are supported in this scenario. You cannot use DTS to migrate incremental data from a standalone MongoDB database because the oplog feature cannot be enabled for the database.</li> <li>DTS cannot migrate data from the admin or local database.</li> <li>Transaction information is not retained. When transactions are migrated to the destination database, they are converted into a single record.</li> <li>To ensure compatibility, the version of the destination MongoDB database must be the same as or later than the version of the source MongoDB database. If the version of the destination database is earlier than the version of the source database, database compatibility issues may occur.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination database. After full data migration is completed, the storage space for collections of the destination database. If source database.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the collections of the destination database. DTS uses the ROUND (COLUMN, PRECISI ON) function to retrieve values from columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN, PRECISI ON) function to retrieve values from columns of the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance after the task is resumed.</li> </ul>
Special cases	If the source database is a self-managed MongoDB database, we recommend that you do not perform a primary/secondary switchover on the database when the data migration task is running. Otherwise, the task fails.

#### Migrate data from a MongoDB database (replica set architecture) to another MongoDB database (replica set architecture or sharded cluster architecture)

Category

Category	Description
Limits on the source database	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select collections as the objects to be migrated and you need to edit collections (such as renaming collections), up to 1,000 collections can be migrated in a single data migration task. If you run a task to migrate more than 1,000 collections, a request error occurs. In this case, we recommend that you split the collections to be migrated, configure multiple tasks to migrate the collections, or configure a task to migrate the entire database.</li> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> <li>The oplog is enabled. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> <li>If you perform only incremental data migration, the oplogs of the source database must be stored for are than 24 hours. If you perform both full data migration and incremental data migration is completed, you can set the retention period to more than 24 hours. Otherwise, DTS may fail to obtain the oplogs and the task may fail. In extreme cases, data may be inconsistent or lost. Make sure that you set the retention period of oplogs based on the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not guarantee service reliability and performance.</li> <li>Limits on operations:</li> <li>During schema migration and full data migration, do not change the schemas of databases or collections. Otherwise, the data migration task fails.</li> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data will be inconsistent between the source and destination databases. To ensure data c</li></ul>

Category	Description
Other limits	<ul> <li>To ensure compatibility, the version of the destination MongoDB database must be the same as or later than the version of the source MongoDB database. If the version of the destination database is earlier than the version of the source database, database compatibility issues may occur.</li> <li>DTS cannot migrate data from the admin or local database.</li> <li>Transaction information is not retained. When transactions are migrated to the destination database, they are converted into a single record.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the collections of the destination database. After full data migration is completed, the storage space for collections of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMM, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the dostination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.</li> </ul>
Special cases	<ul> <li>If the source database is a self-managed MongoDB database, take note of the following limits:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no update operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the migration task is too high, you can perform an update operation on the source database to update the latency.</li> <li>Note If you select an entire database as the object to be migrated, you can create a heartbeat. The heartbeat is updated or receives data every second.</li> </ul>

### 2.11. Precautions and limits for migrating data from a Db2 for LUW database

This topic describes the precautions and limits when you migrate data from a database. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

#### Migrate data from a Db2 for LUW database to a PolarDB-X instance

Category	Description	
	• The server to which the source database belongs must have sufficient outbound bandwidth. Otherwise, the data migration speed decreases.	
	• The tables to migrate must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.	
	• If you select tables as objects to migrate and you need to edit tables (such as renaming tables or columns) in the destination database, up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables and configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.	
	<ul> <li>If you need to migrate incremental data, you must make sure that the following requirements are met:</li> </ul>	
	<ul> <li>The data logging is enabled. Otherwise, error messages are returned during precheck and the data migration task cannot be started.</li> </ul>	
Limits on the source database	<ul> <li>For an incremental data migration, data logs of the source database are retained for at least 24 hours. For a full data and incremental data migration, data logs of the source database are retained for at least seven days. After full data migration is complete, you can set the retention period to more than 24 hours. Otherwise, Data Transmission Service (DTS) may fail to obtain the data logs and the task may fail. In exceptional circumstances, data inconsistency or loss may occur. Make sure that you set the retention period of data logs in accordance with the preceding requirements. Otherwise, the Service Level Agreement (SLA) of DTS does not ensure service reliability and performance.</li> </ul>	
	• Limits on operations:	Limits on operations:
	<ul> <li>During full data migration, do not perform DDL operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>	
		<ul> <li>If you perform only full data migration, do not write data to the source database during data migration. Otherwise, data inconsistency between the source and destination databases may occur. To ensure data consistency, we recommend that you select full data migration and incremental data migration.</li> </ul>

Category	Description
Other limits	<ul> <li>Before you configure a data migration task, you must create databases and tables in the destination instance.</li> <li>DTS migrates incremental updates from a database to the destination database based on the Change Data Capture (CDC) replication technology of . However, the CDC replication technology has its own limits. For more information, see General data restrictions for SQL Replication.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of the tablespace used by the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meet your business requirements. DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the data in the destination instance.</li> </ul>
Special cases	<ul> <li>You must take note of the following items because the source database is a selfmanaged database:</li> <li>If you perform a primary/secondary switchover on the source database when the data migration task is running, the task fails.</li> <li>DTS calculates migration latency based on the timestamp of the latest migrated data in the destination database and the current timestamp in the source database. If no DML operation is performed on the source database for a long time, the migration latency may be inaccurate. If the latency of the migration task is too high, you can perform a DML operation on the source database to update the latency.</li> <li>Note If you select an entire database as the object to migrate, you can create a heartbeat table. The heartbeat table is updated or receives data every second.</li> </ul>

# 2.12. Precautions and limits for migrating data from a Teradata database

This topic describes the precautions and limits when you migrate data from a Teradata database. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

## Migrate data from a Teradata database to an AnalyticDB for PostgreSQL instance

The following table describes the precautions and limits.

Description
<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> </ul>
• If you select tables as the objects to be migrated and you need to edit tables (such as renaming tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.
• Limits on operations:
<ul> <li>During schema migration and full data migration, do not perform data definition language (DDL) operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> </ul>
<ul> <li>In this scenario, DTS does not support incremental data migration. To ensure data consistency, we recommend that you do not write data to the source instance during data migration.</li> </ul>

Category	Description
Category Other limits	<ul> <li>Description</li> <li>You can configure a data migration task for this scenario only in the China (Shanghai), China (Qingdao), or China (Zhangjiakou) region.</li> <li>In this scenario, DTS supports only schema migration and full data migration. DTS does not support incremental data migration.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database is larger than that of the source database.</li> <li>You must make sure that the precision settings for columns of the FLOAT or DOUBLE data type meets your business requirements. DTS uses the ROUND (COLUMN , PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write</li> </ul>
	• DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination instance, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination instance. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

## 2.13. Precautions and limits for migrating data from an HBase database

This topic describes the precautions and limits when you migrate data from a self-managed HBase database. To ensure that your data migration task runs as expected, read the precautions and limits before you configure the task.

## Migrate data from a self-managed HBase database to an AnalyticDB for MySQL cluster V3.0

Category

Category	Description
Limits on the source database	<ul> <li>Bandwidth requirements: The server to which the source database belongs must have sufficient egress bandwidth. Otherwise, the data migration speed is affected.</li> <li>The tables to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.</li> <li>If you select tables as the objects to be migrated and you need to edit tables (such as renaming tables or columns), up to 1,000 tables can be migrated in a single data migration task. If you run a task to migrate more than 1,000 tables, a request error occurs. In this case, we recommend that you split the tables to be migrated, configure multiple tasks to migrate the tables, or configure a task to migrate the entire database.</li> <li>Limits on operations:</li> <li>During schema migration and full data migration, do not perform data definition language (DDL) operations to change the schemas of databases or tables. Otherwise, the data migration task fails.</li> <li>In this scenario, DTS does not support incremental data migration. To ensure data consistency, we recommend that you do not write data to the source instance during data migration.</li> </ul>
Other limits	<ul> <li>In this scenario, DTS supports only schema migration and full data migration. DTS does not support incremental data migration.</li> <li>The characters in the source objects can only be alphanumeric (a to z, A to Z, and 0 to 9). If an object contains other types of characters, schema migration fails.</li> <li>has limits on the usage of disk space. If the disk space usage of the nodes in an cluster reaches 80%, the performance of data writing to the destination database is compromised and the DTS task is delayed. If the usage reaches 90%, data cannot be written to the destination database and error messages are returned. We recommend that you estimate the required disk space based on the objects that you want to migrate. You must make sure that the destination cluster has sufficient storage space.</li> <li>Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers.</li> <li>During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.</li> <li>DTS attempts to resume data migration tasks that failed within the last seven days. Before you switch workloads to the destination cluster, stop or release the data migration task. You can also run the revoke command to revoke the write permissions from the accounts that are used by DTS to access the destination cluster. Otherwise, the data in the source database will overwrite the data in the destination cluster after the task is resumed.</li> </ul>
## 3.Prepare the database accounts for data migration

When you configure a data migration task, you must specify the accounts of the source and destination databases. The database accounts are used for data migration. Different databases and migration types require different permissions. You must create and authorize database accounts before you configure a data migration task.

#### Permissions required for the source database account

Database	Required permission	References	
ApsaraDB RDS for MySQL instance	Read permissions on the objects to migrate	Create databases and accounts for an ApsaraDB RDS for MySQL instance <b>and</b> Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance	
Self-managed MySQL database	<ul> <li>Schema migration: the SELECT permission on the objects to migrate</li> <li>Full data migration: the SELECT permission on the objects to migrate</li> <li>Incremental data migration: the SELECT permission on the objects to migrate, the REPLICAT ION CLIENT, REPLICAT ION SLAVE, and SHOW VIEW permissions, and the permissions to create databases and tables. The permissions allow DTS to create a database named dts to record heartbeat data during migration.</li> </ul>	Create an account for a user-created MySQL database and configure binary logging	
PolarDB for MySQL cluster	Read permissions on the objects to migrate	Create a database account	
PolarDB for Oracle cluster	Permissions of a privileged account	Create database accounts	
ApsaraDB RDS forMariaDB TXnstance		For more information, see Create a database and account on an ApsaraDB RDS for MariaDB TX instance.	

Database	Required permission	References
ApsaraDB RDS for SQL Server instance	<ul> <li>Schema migration: the SELECT permission on the objects to migrate</li> <li>Full data migration: the SELECT permission on the objects to migrate</li> <li>Incremental data migration: the owner permission on the object to migrate</li> <li>Note A privileged account has the required permissions.</li> </ul>	For more information, see Create an account for an ApsaraDB RDS for SQL Server instance.
Self-managed SQL Server database	<ul> <li>Schema migration: the SELECT permission on the objects to migrate</li> <li>Full data migration: the SELECT permission on the objects to migrate</li> <li>Incremental data migration: the permissions of the sysadmin role</li> </ul>	For more information, see CREATE USER.
ApsaraDB RDS for PostgreSQL instance	<ul> <li>Schema migration: the USAGE permission on pg_catalog</li> <li>Full data migration: the SELECT permission on the objects to migrate</li> <li>Incremental data migration: the permissions of a privileged account. The account must be the owner of the database.</li> <li>Note If the source database runs on an ApsaraDB RDS for PostgreSQL 9.4 instance and you migrate only DML operations, the database account must have the REPLICATION permission.</li> </ul>	Create an account on an ApsaraDB RDS for PostgreSQL instance and Create a database on an ApsaraDB RDS for PostgreSQL instance
Self-managed PostgreSQL database	<ul> <li>Schema migration: the USAGE permission on pg_catalog</li> <li>Full data migration: the SELECT permission on the objects to migrate</li> <li>Incremental data migration: permissions of the superuser role</li> </ul>	For more information, see CREATE USER and GRANT.

Database	Required permission	References	
Self-managed Oracle database	<ul> <li>Schema migration: permissions of the schema owner</li> <li>Full data migration: permissions of the schema owner</li> <li>Incremental data migration: permissions of the database administrator (DBA)</li> </ul>	For more information, see CREATE USER and GRANT. Notice If you want to migrate incremental data from an Oracle database but permissions of the DBA cannot be granted to the database account, you can grant fine-grained permissions to the account. For more information, see Migrate data from a self- managed Oracle database to an AnalyticDB for PostgreSQL instance.	
ApsaraDB for MongoDB instance	<ul> <li>Full data migration: read permissions on the source database</li> <li>Incremental data migration: read permissions on the source, admin, and local databases</li> </ul>	For more information, see Manage user permissions on MongoDB databases.	
Self-managed MongoDB database	<ul> <li>Full data migration: read permissions on the source database</li> <li>Incremental data migration: read permissions on the source, admin, and local databases</li> </ul>	For more information, see db.createUser().	
ApsaraDB for Redis instance	Read permissions on the objects to migrate	Create and manage database accounts	
Self-managed Redis database	The PSYNC or SYNC statement can be executed on the source Redis database.	None	
Self-managed TiDB database	The SELECT permission on the objects to migrate and the SHOW VIEW permission	For more information, see Privilege Management	
Self-managed Db2 dat <i>a</i> base	<ul> <li>Schema migration: the SELECT permission on the objects to migrate and the CONNECT permission</li> <li>Full data migration: the SELECT permission on the objects to migrate and the CONNECT permission</li> <li>Incremental data migration: the DBADM authority</li> </ul>	For more information, see Creating group and user IDs for a Db2 database installation and Authorities overview.	

Permissions required for the destination database account

Database	Required permission	Authorization method	
ApsaraDB RDS for MySQL instance	Read and write permissions on the destination database	For more information, see Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account for an ApsaraDB RDS for MySQL instance.	
PolarDB for MySQL cluster	Read and write permissions on the destination database	For more information, see Create database accounts.	
Self-managed MySQL database	The ALL permission on the destination database	For more information, see Create an account for a user-created MySQL database and configure binary logging.	
AnalyticDB for MySQL cluster	<ul> <li>Version 2.0: DTS automatically creates a database account and grants permissions to the account. You do not need to specify the database account.</li> <li>Version 3.0: The read and write permissions are required.</li> </ul>	Version 3.0: Create a database account.	
PolarDB-X 1.0 instance	Read and write permissions on the destination database	For more information, see Manage database accounts.	
ApsaraDB RDS for MariaDB TX instance	Read and write permissions on the destination database	For more information, see Create an account on an ApsaraDB RDS for MariaDB TX instance.	
ApsaraDB RDS for SQL Server instance	Read and write permissions on the destination database	For more information, see Create an account on an ApsaraDB RDS for SQL Server instance.	
Self-managed SQL Server database	The ALL permission on the destination database	For more information, see CREATE USER.	
ApsaraDB RDS for PostgreSQL instance	<ul> <li>Schema migration: the CREATE and USAGE permissions on the migrated objects</li> <li>Full data migration: permissions of the schema owner</li> <li>Incremental data migration: the permissions of the schema owner</li> </ul>	For more information, see Create an account on an ApsaraDB RDS for PostgreSQL instance.	
Self-managed PostgreSQL database	The ALL permission on the destination database	For more information, see CREATE USER and GRANT.	
PolarDB for Oracle cluster	Permissions of the schema owner	For more information, see Create database accounts	

Database	Required permission	Authorization method
Self-managed Oracle database	Permissions of the schema owner	For more information, see CREATE USER and GRANT.
ApsaraDB for MongoDB instance	The dbAdminAnyDatabase permission, the read and write permissions on the destination database, and the read permissions on the local database	For more information, see Manage user permissions on MongoDB databases.
Self-managed MongoDB database	Read and write permissions on the destination database and the read permissions on the local database	For more information, see db.createUser().
ApsaraDB for	If you use the instance password, no authorization is required.	None
Redis instance	If you use a custom account, the read and write permissions are required.	For more information, see Create and manage database accounts.
Self-managed Redis database	The database password must be valid.	None

### **4.Migration task management** 4.1. Object name mapping

## Data Transmission Service (DTS) provides the object name mapping feature. You can use this feature to change the names of one or more objects that are migrated to the destination instance. This topic describes how to use the object name mapping feature when you configure a data migration task.

#### Limits

• You can use the object name mapping feature only when a data migration task is configured and the current step is **Configure Migration Types and Objects**.

Onte Do not use the object name mapping feature after a data migration task is started. Otherwise, data may fail to be migrated.

• MongoDB supports only database name mapping and table name mapping.

#### Configure name mapping for a single database, table, and column

1. In the **Configure Migration Types and Objects** step, move the required objects to the **Selected** section, move the pointer over a database or table, and then click **Edit**.

○ Notice You can use the object name mapping feature only on objects that are supported by DTS.

1.Configure Source and Destination 2.Configure Migration Types and Obj	jects 3.Map name modification A.Precheck
<ul> <li>Migration Types: Schema Migration Full Data Migration</li> <li>Data migration applies to short-term migration scenarios. Typical scenarios between Apsara Stack databases.</li> <li>For long-term data synchronization in real time, use the data synchronization</li> </ul>	✓ Incremental Data Migration i include migrating data to the cloud, scaling and sharding databases, and migrating data on feature.
Available If you search globally, please expand the Q dtstestdata Fables For Yiews Functions Functions For Sys	Selected (To edit an object name or its filter, hover over the object and click Edit.) Learn more. d dtsetstdata (20bjects) customer order Cut
Select All	Remove All
<ul> <li>Name batch change:          <ul> <li>No</li> <li>Yes</li> </ul> </li> <li>Information:         <ul> <li>Data migration only copies the data and schema in the source database and in the source database.</li> <li>DDL operations are not supported during data migration because this can c</li> </ul> </li> </ul>	id saves the copy in the destination database. The process does not affect any data or schema cause migration failures.
	Cancel Previous Save Precheck

#### 2. In the dialog box that appears, specify a name for the object in the destination instance.

? Note

- If the source and destination databases are user-created MySQL databases, ApsaraDB RDS for MySQL, or Apsara PolarDB for MySQL, you can specify the data definition language (DDL) and data manipulation language (DML) statements that you want to migrate.
- If different statements are selected for database name mapping and table name mapping, the statements that are selected for table name mapping prevail.

#### • Database name mapping

In the **Edit Database Name** dialog box, enter the database name that you want to use in the destination instance.

Information: After y	ou edit the source database nam	ne, the name of the destination databa	se is
also upuateu.		Source Database Name:dts	testdat
* DatabaseName:	dtstestdatanew		
DML and DDL			
Statement Filtering Se	elect DDL or DML statements.		
DML Filter:	🖌 insert 🔽 update 🔽 delete	0	
Table: 🔽	create 🔽 alter 🔽 drop 🔽 rei	name 🔽 truncate	
View: 🔽	create 🔽 alter 🔽 drop		
Procedure: 🔽	create 🔽 alter 🔽 drop		
Trig&Func: 🔽	create trigger 🔽 drop trigger 🛽	$\checkmark$ create function $\checkmark$ drop function	
Index: 🔽	create 🔽 drop		

• Table name mapping

In the **Edit Table** dialog box, enter the table name that you want to use in the destination instance.

Edit Table		$\times$
Information: Af corresponding tab	ter you edit the table or column name in the source database, the ble or column nam Source Table Name:customer	
* Table Name:	customernew	
Filter:	DTS supports the WHERE clause in SQL statements. Only data that meets the WHERE clause can be migrated to the destination	
DML and DD	L	
Filtering	<b>Select DDL or DML statements.</b> (1)	
DML Filte	er: 🔽 insert 🔽 update 🔽 delete	
Table:	🗸 create 🗸 alter 🔽 drop 🗸 rename 🗹 truncate	
View:	🗸 create 🔽 alter 🔽 drop	
Procedure:	✓ create ✓ alter ✓ drop	
Trig&Func:	$\checkmark$ create trigger $\checkmark$ drop trigger $\checkmark$ create function $\checkmark$ drop function	
Index:	✓ create ✓ drop	

#### • Column name mapping

In the Edit Table dialog box, enter a new name for each column.

Select All	Column Name	Source Column Nam	le:address
	addressnew		varchar(32)
	id		int(11)
✓	name		varchar(32)
			ОК
? Note	In this step, you can clear	the columns that	do not need to be migrated

- 3. Click OK.
- 4. Configure other parameters that are required for the data migration task.

#### Configure name mapping for multiple tables and columns at a time

1. In the **Configure Migration Types and Objects** step, move the required objects to the **Selected** section.

○ Notice If the last selected object is a database, you cannot configure name mapping for tables or columns by specifying the Name batch change parameter.

2. In the lower part of the page, set the Name batch change parameter to Yes, and then click Advanced Settings.

nfigure Source and Destination Databases 2.Configure Migration Types and Objects	3.Advanced Settings > 4.Precheck	
<ul> <li>Migration Types: Schema Migration</li> <li>Full Data Migration</li> <li>During full data migration, data updates in the source database are not</li> <li>For data consistency, we recommend that you select Schema Migration</li> </ul>	Incremental Data Migration migrated to the destination instance. Full Data Migration, and Incremental Data Migration.	
Available  If you search globally, please expand  If wou search globally, please expand  I w apitestdata  I w custm_info  I w dts I w	Selected (To edit an object name or its filter, hover over the object and click Edit.) Learn more. C d d dstestdata (20bjects) c customer f order	
Select All	Remove All	
<ul> <li>Change:</li> <li>Information:</li> <li>1. Data migration only copies the data and schema in the source database data or schema in the source database.</li> <li>2. DDL operations are not supported during data migration because this operations are not supported during data migrations are n</li></ul>	and saves the copy in the destination database. The process does not affect any an cause migration failures.	Prech

- 3. In the Advanced Settings step, configure name mapping for tables and columns based on your needs.
  - i. Select a range. In this example, select **Section**.

Notice If you select Section, you can search for specific databases and tables by entering their names.

- ii. Select a rule from the drop-down list and configure the rule.
- iii. Click Add Modification Rules.

Notice You can repeat steps i and ii to configure different rules for different tables.

iv. Click Preview to check whether the rules meet your expectations.

	Source instance library Source	instance table I		
ŀ	A total of 1 libraries 2 tables that m	eet the above conditions		
Selection Rules:	Change Name   Change	Name: Name		
	Add Modification Rules Pr	review		
dify Rules1:	Table Name(customer)Add Prefix:	test_,Suffix: X Modify Rules2: Table I	Name(order)Change Name: neworder X	
dify Rules1:	Table Name(customer)Add Prefix:	test_Suffix: X Modify Rules2: Table I	Name(order)Change Name: neworder X	
dify Rules1: erall Previev	Table Name(customer)Add Prefix:	test_Suffix: X Modify Rules2: Table I	Name(order)Change Name: neworder X	
dify Rules1: rerall Previev	Table Name(customer)Add Prefix:	test_Suffix: X Modify Rules2: Table I	Name(order)Change Name: neworder ×	
dify Rules1: erall Previev ource Datab	Table Name(customer)Add Prefix: Partial Preview Dase	test_Suffix: X Modify Rules2: Table I Destination Database(Click on the	Name(order)Change Name: neworder ×	
dify Rules1: rerall Previev ource Datab	Table Name(customer)Add Prefix:      Partial Preview  base  fata (2)	test_Suffix: X Modify Rules2: Table I Destination Database(Click on the A dtstestdata	Name(order)Change Name: neworder × name below to edit directly) Batch Change Table Name	
dify Rules1: erall Preview ource Datab dtstestd	: Table Name(customer)Add Prefix: w Partial Preview base fata (2) mer	test_Suffix: X Modify Rules2: Table	Name(order)Change Name: neworder × name below to edit directly) Batch Change Table Name Column Name Modification	
dify Rules1: rerall Preview ource Datab	: Table Name(customer)Add Prefix: w Partial Preview base Jata (2)	test_Suffix: X Modify Rules2: Table I Destination Database(Click on the dtstestdata	Name(order)Change Name: neworder × name below to edit directly) Batch Change Table Name	
dify Rules1 : rerall Previev ource Datab dtstestd custon custon	: Table Name(customer)Add Prefix: w Partial Preview base fata (2) mer	test_Suffix: X Modify Rules2: Table  Destination Database(Click on the  dtstestdata  test_customer  neworder	Name(order)Change Name: neworder × name below to edit directly) Batch Change Table Name Column Name Modification Column Name Modification	

- v. On the **Partial Preview** tab, click the view icon next to the database name in the **Destination Database** section, and then click **Column Name Modification**.
- vi. Configure mapping rules for column names, and then click OK.

Column selection and colur	nn name ch	anges		×
customer table 3 columns				
Filter:	1		Verify	
Source instance colum	Search		2	3
Name Added Suffix 🔻	Prefix: test_	Suffiz Su	к: ıffix	ОК
Modify Rules 1: Table Nar	mecustomer	Column NameAdd Pr	efix: test_,Suffix:	×
Original list name Ty	pe	Target column name	е	
🔝 address 🛛 va	rchar(32)	test_address		
📰 id 🛛 int	(11)	test_id		
🔝 name va	rchar(32)	test_name		
			Cance	el OK

#### ♦ Notice

- By default, a mapping rule applies to all column names. You can search for a specific column name and modify the mapping rule. To do this, you can enter the column name in the search box, and then click Search.
- In the dialog box, you can specify conditions to filter data. For more information, see Use conditions to filter data.

#### vii. Click OK.

- 4. Click Precheck.
- 5. Configure other parameters that are required for the data migration task.

# 4.2. Specify the capitalization of object names in the destination instance

When you configure a Data Transmission Service (DTS) task, you can specify the capitalization of database names, table names, and column names in the destination instance.

#### Limits

When you specify the capitalization of object names in the destination instance, take note of the following limits:

- You cannot specify capitalization rules for views, functions, or stored procedures. If a table in the destination database is renamed, the corresponding views, functions, and stored procedures may fail to be created.
- The objects cannot contain CHECK constraints or computed columns.

#### Capitalization rules

DTS supports the following capitalization rules for object names:

• DTS default policy

DTS configures the capitalization of database names, table names, and column names in the destination instance based on the database type and related parameters, such as <code>lower\_case\_tabl</code> <code>e names</code>.

Destination database type	Destination database parameter	Capitalization of object names in the destination instance	
<ul> <li>Self-managed MySQL database or ApsaraDB RDS for MySQL instance</li> <li>ApsaraDB RDS for MariaDB TX</li> <li>PolarDB for MySQL</li> <li>PolarDB-X</li> <li>AnalyticDB for MySQL V3.0</li> </ul>	The lower_case_table_names parameter is set to 1 or 2.	Database names and table names in the destination instance are in lowercase. The capitalization of column names in the destination instance is the same as that in the source instance.	
	The lower_case_table_names parameter is set to 0.	The capitalization of database names, table names, and column names in the destination instance is the same as that in the source instance.	
<ul> <li>AnalyticDB for MySQL V2.0</li> <li>DataHub</li> <li>MaxCompute</li> <li>HybridDB for MySQL</li> </ul>	None	Database names, table names, and column names in the destination instance are in lowercase.	
Self-managed Oracle database	None	Database names, table names, and column names in the destination instance are in uppercase.	

Destination database type	Destination database parameter	Capitalization of object names in the destination instance
<ul> <li>Self-managed SQL Server database or ApsaraDB RDS for SQL Server instance</li> <li>PolarDB O Edition</li> <li>Self-managed PostgreSQL database or ApsaraDB RDS for PostgreSQL instance</li> <li>ApsaraDB RDS for PPAS</li> <li>AnalyticDB for PostgreSQL</li> <li>Self-managed Db2 database</li> <li>Self-managed MongoDB database or ApsaraDB for MongoDB instance</li> <li>Self-managed Redis database or ApsaraDB for Redis instance</li> <li>Tablestore</li> <li>Elasticsearch</li> </ul>	None	Database names and table names in the destination instance are in lowercase. The capitalization of column names in the destination instance is the same as that in the source instance.

- Consistent with the source database The capitalization of database names, table names, and column names in the destination instance is the same as that in the source instance.
- Consistent with the default policy of the destination database (uppercase) The names of all the databases, tables, and columns that are migrated or synchronized to the destination instance are in uppercase.
- Consistent with the default policy of the destination database (lowercase) The names of all the databases, tables, and columns that are migrated or synchronized to the destination instance are in lowercase.

### 4.3. Filter the data to be migrated

When you configure the objects to be migrated in a data migration task, you can specify conditions to filter data. Only the data that meets the specified conditions is migrated to the destination database. This feature is applicable to scenarios such as regular data migration and table partitioning.

#### Procedure

1. In the **Configure Migration Types and Objects** step, move the required objects to the **Selected** section, move the pointer over a database or table, and then click **Edit**.

<ul> <li>Migration Types: Schema Migration Full Data Mig Data migration applies to short-term migration scenarios. Typic between Apsara Stack databases.</li> <li>For long-term data synchronization in real time, use the data s</li> </ul>	ration Incremental I cal scenarios include migratin ynchronization feature.	Data Migration	tabases, and migrating data
Available  If you search globally, please expand the Q  C C C C C C C C C C C C C C C C C C	> <	Selected (To edit an object name or its f Edit.) Learn more.	ilter, hover over the object and c
Select All		Remove All	
*Name batch change:	database and saves the copy	in the destination database. The process do	pes not affect any data or schema

- 2. In the Edit Table dialog box, enter a condition in the Filter field.
  - Relational dat abases, such as MySQL and SQL Server.

#### Dat a Migrat ion Migrat ion t ask mana gement

Edit Table	5		×
Informa correspo	ation: Af	ter you edit the table or column name in the s le or column name in the destination database	ource database, the e is also updated.
* Table	e Name:	customer	
	Filter:	orderid>100	Verify
✓ Selen All	ct Colu	umn Name	Туре
<b>&gt;</b>	ad	dress	varchar(32)
<b>&gt;</b>	or	derid	int(11)
<b>&gt;</b>	na	me	varchar(32)
			ОК

#### ? Note

- An SQL condition is a standard SQL WHERE statement. The following operators are supported:
   , != , < , and > . Only the data that meets the WHERE condition is migrated to the destination database. In this example, enter orderid>10
   .
- You can specify a time condition in an SQL WHERE statement. However, you must make sure that the specified time condition is valid. For example, to filter incremental data created after 2020, you must enter create\_time>'2020-01-01' or create\_ti me>'2020-01-01 00:00:00'. You cannot enter create time>'2020'.
- You can use apostrophes (') in a filter condition if necessary. For example, you can enter address in ('hangzhou', 'shanghai')
- Filter conditions are case-insensitive. If a table in the source database contains
   Column A and Column a, you can use an SQL WHERE statement to filter only Column A.
   In this case, you can enter where A=10 rather than where a=10.

#### • Non-relational dat abases, such as MongoDB.

Edit Collection Na	me	$\times$
Information: If collection name in	a collection name in the source database is edited, the corresponding the destination database is also updated.	
* Collection Name:	customer	
Filter:	{"id":{\$gt:"52b64e55dc4449b7ba3d1183a0ea171b"}}	
	ок	

Note You can specify a JSON-formatted condition. Only the data that meets the specified condition is migrated to the destination database. In this example, enter {"id":{\$ gt:"52b64e55dc4449b7ba3d1183a0ea171b"}}. Only the data with an ID greater than this value is migrated.

3. Click **Verify** to check whether the syntax is valid.

#### ? Note

- If the syntax is valid, the Information message shows that the validation is passed.
- If the syntax is invalid, the **Error** message appears. In this case, you must modify the condition based on the instructions.
- 4. Click OK.
- 5. Configure other parameters that are required for the data migration task.

## 4.4. View the progress of a data migration task

This topic describes how to view the progress of a data migration task in the DTS console. DTS provides the following migration types: schema migration, full data migration, and incremental data migration.

#### Prerequisites

A data migration task is started.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the data migration instance resides.

- 4. On the Migration Tasks page, click the ID of the data migration instance.
- 5. Perform operations based on the migration type.

**Note** If you do not select a migration type when configuring the task, or the migration type is not supported, this migration type will not be displayed in the left-side navigation pane.

- View the details of schema migration.
  - a. In the left-side navigation pane, choose **Migration Details > Schema Migration**.
  - b. Click the tab that indicates a type of schema object to view the migration details. For example, click **Tables** to view the migration details of tables. You can search for the migration details of a specific object by specifying the object name.

Configure Task	Searc	Search: Search by object name. Search										
<ul> <li>Migration Details</li> </ul>												
Schema Migration	Ta	ables 2	Views	Sequences	Functions	User Defined Types	Rule	Domains	Operators	Aggregates	Extensio	ns
Full Data Migration	Obj	ject Name			Source	Database Desti	ation Datab	ase Status				Actions
Incremental Data Mig	ord	ler			testsche	ema testso	hema	Comple	ted	View CREA	ATE Syntax	View Foreign Key of Index☆
Performance Monitori	0	)bject			Туре			Status				Actions
Configure Monitoring	ti	bl_user_pkey			Constra	int		Comple	ted		View CF	REATE Syntax
Task Log	cus	stomer			testsche	ema testso	hema	Comple	ted	View CREA	ATE Syntax	View Foreign Key of Index⊗
								Total: 2 iten	n(s), Per Page	e: 20 item(s)	« «	1 > »

**Note** In the Actions column of an object, click View Creation Syntax or choose View Index Foreign Keys > View Creation Syntax. In the View Creation Syntax message that appears, you can view the creation syntax of the object.

/iew CREATE Syntax	×
CREATE TABLE "testschema"."customer" ( "id" integer NOT NULL, "user_id" bigint, "user_name" character varying(64), "create_time" timestamp(6) with time zone DEFAULT clock_timestamp() );	
	Close

- View the details of full data migration
  - a. In the left-side navigation pane, choose Migration Details > Full Data Migration.

b. View the migration details of each object or search for the migration details of a specific object by specifying the object name.

Configure Task	Search	n: Search b	y object name.	Searc	ch				
<ul> <li>Migration Details</li> </ul>									
Schema Migration	Obje	ct Name	Source Database	Destination Database	Estimated Keys	Completed	Time Consumed (Seconds)	Status	Actions
Full Data Migration	orde	r	testschema	testschema	1000000	1000000	20Minutes36Seconds	Completed	
Full Data Migration	custo	omer	testschema	testschema	10000	10000	43Seconds	Completed	
Incremental Data Mig									
Performance Monitori						Total: 2	e item(s), Per Page: 20 item(s)	« < 1	> >>
Configure Monitoring									
Task Log									

- View the details of incremental data migration
  - a. In the left-side navigation pane, choose Migration Details > Incremental Data Migration.
  - b. View the migration details of each object or search for the migration details of a specific object by specifying the object name.

Configure Task	Search: Search by o	Search by object name. Search					
<ul> <li>Migration Details</li> </ul>							
Schema Migration	Object Name	Source Database	Destination Database	Status (All)			
Full Data Migration	order	testschema.testschema	testschema.testschema	The migration task is not delayed.			
Incremental Data Mig	customer	testschema.testschema	testschema.testschema	The migration task is not delayed.			
Performance Monitori			т	fotal: 2 item(s), Per Page: 20 item(s) $\begin{tabular}{ c c c c c } & & & & \\ & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & $			
Configure Monitoring							
Task Log							

#### Reference

- View the connection state and performance of full data migration: View the bandwidth, records per second (RPS), read/write response time, and network latency.
- View the connection status and performance of incremental data migration: View the number of migrated rows, bandwidth, and migration performance.

## 4.5. View the connection state and performance of full data migration

This topic describes how to view the connection state and performance of full data migration in the Data Transmission Service (DTS) console. DTS provides the following connection and performance metrics: bytes per second (BPS), records per second (RPS), read/write response time, and network latency. You can monitor and manage data migration tasks by using these metrics.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the data migration instance resides.
- 4. On the Migration Tasks page, click the ID of the data migration instance.
- 5. In the left-side navigation pane, choose **Performance Monitoring > Performance of Full Data**

#### Migration.

6. On the page that appears, the connection state and performance of full data migration are displayed. You can select a time range to view the trend charts of performance metrics for full data migration.

Full migration lin	k topology 🏾 🕽		Adjust full migration rate
Sou	Network Delay: 0.0 ms BPS: 5.10 MB/s RPS: 6389.00 Row/s		Network Delay: 0.0 ms Write RT: 1 ms BPS: 5:39 MB/s RPS: 6747.00 Row/s
Full Data Migratio	on Performance		
Select Range:	Jul 1, 2020, 15:13 - Jul 1, 2020, 16:13		
		Indicator Meaning	Indicator Meaning
	Bandwidth (Mbit/s)		Source, Target Instance RPS
20		25,000	
		20,000	
10		15,000	
		10,000	

Section	Description
	In this section, you can view the read/write performance and network information about the connections between DTS and the source and destination databases. The following parameters are provided:
	<ul> <li>Connection between DTS and the source database</li> </ul>
	<ul> <li>BPS: the amount of data that DTS reads from the source database per second. Unit: MB/s.</li> </ul>
	<ul> <li>RPS: the number of records that DTS reads from the source database per second.</li> </ul>
Topology of Full	<ul> <li>Network Latency: the network latency between DTS and the source database.</li> </ul>
Data Migration	• Connection between DTS and the destination database
	<ul> <li>BPS: the amount of data that DTS writes to the destination database per second. Unit: MB/s.</li> </ul>
	<ul> <li>RPS: the number of records that DTS writes to the destination database per second.</li> </ul>
	<ul> <li>Network Latency: the network latency between DTS and the destination database.</li> </ul>
	<ul> <li>Write RT: the response time period when DTS writes data to the destination database.</li> </ul>

Section	Description				
Performance of Full Data Migration	In this section, you can view the bandwidth, RPS, read/write response time, and network latency.				
	<b>Note</b> To view the description of performance metrics, move the pointer over the <b>Indicator Meaning</b> button at the upper-right corner of a trend chart.				

# 4.6. View the connection status and performance of incremental data migration

This topic describes how to view the connection status and performance of incremental data migration in the Data Transmission Service (DTS) console. DTS provides the following connection and performance metrics: the number of DDL operations, number of slow SQL statements in the destination database, bandwidth, migration performance, and migration latency. You can monitor and manage data migration tasks by using these metrics. In addition, DTS provides the task diagnostics feature.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the data migration instance resides.
- 4. On the Migration Tasks page, click the ID of the data migration instance.
- 5. In the left-side navigation pane, choose **Performance Monitoring > Performance of Incremental Data Migration**.
- 6. On the page that appears, view the topology and diagnostic details of the migration task. You can also select a time range to view the trend charts of performance metrics.



Section	Description	
Topology of Incremental Data Migration	<ul> <li>In this section, you can view the data transmission status and network latency between DTS modules. The DTS modules include the source database, data collection module, data caching module, data writing module, and destination database. The following parameters are provided:</li> <li>BPS: the bandwidth between DTS modules. Unit: MB/s.</li> <li>RPS: the number of records that are transmitted between DTS modules per second.</li> <li>Network Latency: the network latency between DTS modules.</li> </ul>	
Diagnostics	You can check the performance of the source database, destination database, network, and DTS during incremental data migration. DTS provides diagnostic results and suggestions. To use this feature, perform the following steps: <ul> <li>i. In the upper-right corner of the page, click <b>Diagnostics</b>.</li> <li>ii. In the message that appears, click <b>OK</b>.</li> <li>iii. Refresh the page to view the diagnostic progress. To view diagnostic details, click the following figure.</li> </ul>	
Performance of Incremental Data Migration	<ul> <li>In this section, you can view the bandwidth, migration performance, and migration latency. You can also select other metrics from the More Metrics drop-down list.</li> <li>The following metrics are provided:</li> <li>Bandwidth: the amount of data that is written to the destination datab per second. Unit: MB/s.</li> <li>Migration Performance (RPS): the number of records that are written the destination database per second.</li> <li>Migration Latency (s): the difference between the timestamp of the latest migrated data in the destination database and the current timesta in the source database.</li> <li>Number of DDL Operations: the number of DDL operations that are executed in the destination database and the time when a data record is generated in the source database. For example, if a data record is generated in the source database at 7 o'clock and DTS writes the data record to the destination database at 8 o'clock, the task is delayed by 1 hour.</li> <li>Slow SQL Queries: the number of slow SQL queries that are generated the database within the selected time range.</li> </ul>	

## 4.7. Modify the transfer rate of full data migration

This topic describes how to modify the transfer rate of full data migration in the Data Transmission Service (DTS) console.

#### Limits

To modify the transfer rate of full data migration, you must make sure that the source and destination databases belong to the types listed in the following table.

Source database	Destination database
<ul><li>Self-managed MySQL database</li><li>ApsaraDB RDS for MySQL</li></ul>	<ul><li>Self-managed MySQL database</li><li>ApsaraDB RDS for MySQL</li><li>MaxCompute</li></ul>
<ul><li>Self-managed PostgreSQL database</li><li>ApsaraDB RDS for PostgreSQL</li></ul>	<ul><li>Self-managed PostgreSQL database</li><li>ApsaraDB RDS for PostgreSQL</li></ul>
Self-managed Oracle database	<ul> <li>Self-managed Oracle database</li> <li>Self-managed MySQL database</li> <li>ApsaraDB RDS for MySQL</li> <li>Self-managed PostgreSQL database</li> <li>ApsaraDB RDS for PostgreSQL</li> <li>PolarDB for Oracle</li> <li>AnalyticDB for MySQL (version 3.0)</li> <li>AnalyticDB for PostgreSQL</li> </ul>
PolarDB for Oracle	PolarDB for Oracle
Self-managed TiDB database	<ul><li>Self-managed MySQL database</li><li>ApsaraDB RDS for MySQL</li><li>PolarDB for MySQL</li></ul>

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the data migration instance resides.
- 4. On the Migration Tasks page, click the ID of the data migration instance.
- 5. In the left-side navigation pane, choose **Performance Monitoring > performance of Full Data Migration**.
- 6. In the upper-right corner of the page, click **Adjust full migration rate**.

7. Modify the transfer rate of full data migration based on your needs.

Adjust full migration rate		$\times$
Limit full migration rate? 💿 Yes 🔘	No	
The current full migration rate is		
The per second rate of	Adjusted	
queries to the original library:	to 500	
Rows migrated in full per	Adjusted	
second RPS:	to 2000	
MB of data migrated in full	Adjusted	
per second:	to 20	
	ОК	Close
? Note		
You can click the 🕜 icon to view th	he details of each parameter.	

8. Click OK.

### 4.8. Fix a failed data migration task

This topic describes how to fix a failed data migration task. You can use this feature if your data migration task is in the **Migration Failed** state during schema migration or full data migration.

#### Fix a failed task during schema migration

DTS supports data migration between heterogeneous data sources. However, if you migrate data of unsupported types to the destination instance during schema migration, the task fails.

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the destination instance resides.
- 4. Use one of the following methods to fix the failed task:
  - Method 1
    - a. Find the target task and click View Cause and Rectify.

	Task ID/Name:	Status: Migration Failed View Cause and Rectify Start Task View Details	Duplicate Task   Upgrade   Configure Monitoring and Alerting Completed

- b. Fix the issue based on the cause of failure that is displayed in the View Cause and Rectify message. For example, you can fix an issue by modifying the schema syntax.
- c. Click Restart Task.

```
• Method 2
```

- a. Click the ID of the target task.
- b. In the left-side navigation pane, choose Migration Details > Schema Migration.
- c. On the **Schema Migration** page, find the object that causes the migration failure and click **Rectify** in the Status column.

	<	<b>?</b>	-	10						
	Configure Task	Search: Search	h by object na	ame.		Search				
•	Migration Details		_							
	Schema Migr <b>.(TimRectify)</b>	Tables 1	Views	Functions	Procedures					
	Full Data Migration	Object Name			Source I	Database	Destination Database	e Status		Actions
	Incremental Data Mig	order			dtstestd	ata	dtstestdata	Failed Rectify, Ignore	View CREATE Syntax	View Foreign Key of Index⊗
•	Performance Monitori							T 1 1 4 7 - ( )		
	Configure Monitoring							i otai: 1 item(s), Per Page: 2	(U ITEM(S) « <	

- d. Fix the issue based on the cause of failure that is displayed in the **Rectify** dialog box. For example, you can fix an issue by modifying the schema syntax.
- e. Click **Rect if y**.
  - ? Note
    - If the failure is not fixed, the Rectify dialog box does not close and shows the cause of failure. You must continue troubleshooting based on the cause of failure until the troubleshooting is successful.
    - If the troubleshooting is successful, the Schema Migration page appears and the status of the object changes to Finished.
- 5. If no objects are in the Failed state, DTS proceeds with the data migration task, for example, entering the full data migration process.

#### Fix a failed task during full data migration

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination instance resides.
- 4. Find the target task and click **View Cause and Rectify**. DTS allows you to fix a task that fails during full data migration due to the following reasons.

**?** Note If a task fails during full data migration due to other reasons, DTS provides only the **Ignore** option. The object that causes the failure is not migrated to the destination database.

- The connection to the source or destination database failed or timed out.
   Troubleshoot the issue, make sure that the connection is successful, and then click Restart Task.
- The storage space of the destination instance is insufficient or the instance is locked. Upgrade the specification of the destination instance or clear the log space, and then click **Restart Task**.
- MyISAM tables in the source database are corrupted. Manually fix the issue in the source database, and then click **Restart Task**.

- 5. In the dialog box that appears, fix the issue based on the cause of failure.
- 6. Click Restart Task.

## 4.9. Switch workloads to the destination database

This topic describes how to switch your workloads to the destination database and prepare a rollback solution. This allows you to minimize the negative impact of data migration on your business.

#### Prerequisites

A data migration task is configured and it is in the **Migrating** or **Completed** state. For more information, see Overview of data migration scenarios.

#### Precautions

- We recommend that you switch workloads to the destination database during off-peak hours to minimize the negative impact. Before you switch workloads to the destination database, you must stop writing data to the source database and suspend the business.
- We recommend that you create and authorize a database account for data migration. This allows you to distinguish session information and improve data security.

#### Procedure

1. Wait until the task progress bar shows **Incremental Data Migration** and **The migration task is not delayed** or a delay time of less than 5 seconds.

Task ID/Name:	Status: Migrating Pause Task View Details Duplicate Task Up	grade Configure Monitoring and Alerting Modify password
2019-12-18 16:24:40 Created		Completed
Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 49752)	Incremental Data Migration The migration task is not delay

(?) Note If you do not select Incremental Data Migration when you configure the data migration task, the task progress bar does not show Incremental Data Migration. After data is migrated, the migration task automatically ends. In this case, you must suspend the business and stop writing data to the source database before you run the data migration task. Skip to Step 5 and proceed.

- 2. Suspend the business and stop writing data to the source database.
- 3. Log on to the source database and run the following statements based on the database type to view the session information. Make sure that no new sessions are used for write operations.



4. After the status of **incremental data migration** changes to **The migration task is not delayed** again, wait for one minute or longer, and then manually stop the migration task.

k Name v Search by migration task name.	Search Sort: Default Sorting V Status: All V	Tag
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	CUpgrade   Configure Monitoring and Alerting   Change password   Edit Tag
4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed Incremental Data M 1 n The migration task is not delayed.
3 Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) < < 1 > >

- 5. During service interruption, remove the limit on writing data to the source database.
- 6. Create and start a task to migrate incremental data generated in the destination database to the source database. The migration task created in this step provides a rollback solution. If an error occurs in the destination database, you can switch workloads to the source database.

For example, data is migrated from a self-managed MySQL database to an ApsaraDB RDS for MySQL instance. To create a task in the opposite direction, see Migrate data from an ApsaraDB RDS for MySQL instance to a self-managed MySQL database (select only Incremental Data Migration).

• Warning When you configure a data migration task in the opposite direction, you must select only Incremental Data Migration in the Configure Migration Types and Objects step. Then, you must select the database or table to be migrated back to the source database.

1.Configure Source and Destination 2.Configure Migration Types and Obj	jects	3.Map name modification	>	4.Precheck
<ul> <li>Migration Types: Schema Migration Full Data Migration</li> <li>Data migration applies to short-term migration scenarios. Typical scenario between Apsara Stack databases.</li> <li>For long-term data synchronization in real time, use the data synchronization</li> </ul>	Incremental Dat is include migrating of tion feature.	a Migration	atabases, and mig	irating data
Available  If you search globally, please expand the Q  def data123  mysqltest  Tables  Ge trabes  Ge trabes  For Cedures  Consecutives  Conse	> <	Selected (To edit an object name or its i Edit.) Learn more.	filter, hover over I	the object and click
Select All		Remove All		
*Name batch change:      No      Yes  Information:  Data migration only copies the data and schema in the source database an in the source database.  DDL operations are not supported during data migration because this can	nd saves the copy in cause migration failu	the destination database. The process d	oes not affect any	r data or schema
			Cancel	Previous Save Precheck

- 7. Verify that the data of the source and destination databases is consistent, switch workloads to the destination database, and then resume your business.
- 8. After you run the task in the opposite direction, incremental data generated in the destination database is migrated back to the source database in real time. If the business fails, you can switch workloads back to the source database.

#### What to do next

After you switch workloads to the destination database and test all the business-related features, you can stop the task in the opposite direction. For more information, see Stop a data migration task.

• Warning The database accounts that are used for data migration have the read and write permissions. After data is migrated, you must delete the accounts or revoke the write permission to ensure security.

#### FAQ

- Q: What can I do if an error occurs after I switch workloads to the destination database? A: If an error occurs, you can switch workloads back to the source database. After you run the task in the opposite direction, incremental data generated in the destination database is migrated back to the source database in real time.
- Q: How can I ensure data consistency in the source database if I am unable to switch workloads to the destination database?

A: You can back up the source database before you switch workloads.

Q: What can I do if data is written to the source database due to a misoperation after I switch workloads to the destination database?
 A: You can compare data of the source and destination database through data verification and manually change data to ensure consistency.

### 4.10. Stop a data migration task

This topic describes how to stop a data migration task. If you do not need a task or the task failed, you can stop the task. This ensures that data from the source database does not overwrite data in the destination database.

#### Prerequisites

The data migration task is in the Migrating, Paused, or Migration Failed state.

#### Impact

A stopped data migration task is in the **Completed** state. You can only delete the task when it is in this state.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. Perform operations based on the selected migration type.

• Warning To minimize the negative impact of data migration on your business, we recommend that you switch your workloads to the destination instance and prepare a rollback solution. For more information, see Switch workloads to the destination database. If you do not need to switch your workloads, you can stop the migration task by using the following procedure.

• Full data migration

Do not manually stop a task during full data migration. Otherwise, the system may fail to migrate all data. Wait until the migration task automatically ends.

• Incremental data migration

The task does not automatically end during incremental data migration. You must manually stop the migration task.

- a. Wait until the task progress bar shows **Incremental Data Migration** and **The migration task is not delayed**. Then, stop writing data to the source database for a few minutes. In some cases, the progress bar shows the delay time of **incremental data migration**.
- b. After the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.

Task Name            Search by migration task name.	Search Sort: Default Sorting V Status: All V T	ig
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed
Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) < < 1 > >

### 4.11. Data formats of a Kafka cluster

When you use Data Transmission Service (DTS) to migrate or synchronize data to a Kafka cluster, you can select the format in which data records are stored. This topic describes the data formats that are supported by DTS. You can parse data based on the definition of these data formats.

#### Data formats

DTS allows you to store data written to a Kafka cluster in the following three formats:

- DTS Avro: A data serialization format into which data structures or objects can be converted to facilitate storage and transmission.
- SharePlex JSON: The data replication software SharePlex reads the data in the source database and writes the data to the Kafka cluster in the SharePlex JSON format.
- Canal JSON: Canal parses the incremental logs of the source database and transfers the incremental data to the Kafka cluster in the Canal JSON format.

#### DTS Avro

**DTS Avro** is the default data format. Data migrated or synchronized to a Kafka cluster by DTS is stored in the Avro format. You need to parse data based on the schema definition of DTS Avro. For more information, visit Git Hub.

#### SharePlex JSON

Parameters

Parameter	Description
time	The UTC time when the transaction in the database is committed, in the yyyy-MM-ddTHH:mm:ssZ format.
userid	The ID of the user who commits the transaction.
qo	The operation type. Valid values: INSERT, UPDATE, DELETE, TRUNCATE, DROP COLUMN, UPDATE BEFORE, and UPDATE AFTER.

Parameter	Description
scn	The system change number (SCN) that identifies the version of the transaction that the database commits at a specific time. Each committed transaction is assigned a unique SCN.
rowid	A relatively unique address value that is used to identify a record in the database.
trans	The ID of the transaction.
seq	The sequence number of the operation in the transaction, starting from 1.
size	The total number of operations in the transaction.
table	The name of the table.
idx	The index of the operation in the transaction, in the seq/size format. For example, 1/11 indicates that the sequence number of the operation is 1 in the transaction with a total number of 11 operations.
posttime	The time when the transaction is committed to the destination database.

#### Examples:

#### • Insert data:

```
{
   "meta": {
       "time": "2017-06-16T14:24:34",
       "userid": 84,
       "op": "ins",
         "scn": "14589063118712",
         "rowid": "AAATGpAAIAAItcIAAA",
       "trans": "7.0.411499",
       "seq": 1,
       "size": 11,
       "table": "CL_BIZ1.MIO_LOG",
        "idx": "1/11",
       "posttime": "2017-06-16T14:33:52"
   },
    "data": {
       "MIO_LOG_ID": "32539737"
    }
}
```

• Update data:

#### Dat a Transmission Service

```
{
   "meta": {
       "time": "2017-06-16T15:38:13",
       "userid": 84,
       "op": "upd",
       "table": "CL_BIZ1.MIO_LOG"
       . . .
   },
   "data": {
       "CNTR NO": "1171201606"
   },
   "key": {
       "MIO_LOG_ID": "32537893",
       "PLNMIO_REC_ID": "31557806",
       "POL_CODE": null,
       "CNTR_TYPE": null,
       "CNTR NO": "1171201606syui26"
   }
}
```

#### • Delete data:

```
{
    "meta": {
        "time": "2017-06-16T15:51:35",
        "userid": 84,
        "op": "del",
     },
    "data": {
        "MIO_LOG_ID": "32539739",
        "PLNMIO_REC_ID": "31557806",
        "POL_CODE": null,
        "CNTR_TYPE": null,
        "CG_NO": null
     }
}
```

#### Canal JSON

#### Parameters

Parameter	Description		
database	The name of the database.		
es	The time when the operation is performed in the database. The value is a 13-bit UNIX timestamp. Unit: ms.		
	<b>Note</b> You can use a search engine to obtain a UNIX timestamp converter.		
id	The serial number of the operation.		

Parameter	Description		
isDdl	<ul><li>Indicates whether the operation is a data definition language (DDL) operation.</li><li>true: yes</li><li>false: no</li></ul>		
mysqlType	The data type of the field.		
old	The data before update.		
pkNames	The name of the primary key.		
sql	The SQL statement.		
sqlType	The converted field type. For example, unsigned int is converted to Long, and unsigned long is converted to BigDecimal.		
table	The name of the table.		
ts	The time when the operation is written to the destination database. The value is a 13-bit UNIX timestamp. Unit: ms.		
	<b>ONDE</b> You can use a search engine to obtain a UNIX timestamp converter.		
type	The operation type. Valid values: DELETE, UPDATE, and INSERT.		

The following script shows an example of UPDATE operation:

```
"data": [
 {
    "id": "500000287",
    "shipping_type": null
  }
],
"database": "dbname",
"es": 1600161894000,
"id": 58,
"isDdl": false,
"mysqlType": {
 "id": "bigint(20)",
  "shipping type": "varchar(50)"
},
"old": [
 {
    "shipping type": "aaa"
 }
],
"pkNames": [
 "id"
],
"sql": "",
"sqlType": {
 "id": -5,
 "shipping_type": 12
},
"table": "tablename",
"ts": 1600161894771,
"type": "UPDATE"
```

```
]
```

## 4.12. Specify the policy for migrating data to Kafka partitions

When you configure a task to migrate data to a Kafka cluster, you can specify the policy for migrating data to Kafka partitions. The policy allows you to improve the migration performance. For example, you can migrate data to different partitions based on hash values.

#### Hash algorithm

Data Transmission Service (DTS) uses the hashCode() method in Java to calculate hash values.

#### **Configuration method**

In the **Configure Migration Types and Objects** step of a task creating wizard, you can specify the policy for migrating data to Kafka partitions. For more information, see Migrate data from a self-managed Oracle database to a Message Queue for Apache Kafka instance and Overview of data migration scenarios.

Q Warning After a data migration task is started, do not change the number of partitions in the destination topic. Otherwise, data migration fails.

#### Policies

Policy	Description	Advantage and disadvantage
Ship All Data to Partition 0	DTS migrates all data and DDL statements to Partition 0 of the destination topic.	<ul> <li>Advantage: The order in which all objects are created and changed is the same as that in the source database.</li> <li>Disadvantage: This policy provides ordinary migration performance.</li> </ul>
Ship Data to Separate Partitions Based on Hash Values of Database and Table Names	DTS uses the database and table names as the partition key to calculate the hash value. Then, DTS migrates the data and DDL statements of each table to the corresponding partition of the destination topic.	<ul> <li>Advantage: The order in which a destination table is created and changed is the same as that of the source table. This policy provides good migration performance.</li> <li>Disadvantage: Tables are migrated to different partitions. After data migration, the order of data changes on different tables may become inconsistent.</li> </ul>
	<ul> <li>Note</li> <li>The data and DDL statements of the same table are migrated to the same partition.</li> <li>If a DDL statement is irrelevant to a table, for example, CREAT E DAT ABASE, the statement is migrated to Partition 0.</li> </ul>	
Ship Data to Separate Partitions Based on Hash Values of Primary Keys	DTS uses a table column as the partition key to calculate the hash value. The table column is the primary key by default. If a table does not have a primary key, the unique key is used as the partition key. DTS migrates each row to the corresponding partition of the destination topic. You can specify one or more columns as partition keys to calculate the hash value.	<ul> <li>Advantage: This policy provides the best migration performance.</li> <li>Disadvantage: After data migration, the order of data changes on each data record remains the same. However, the order of data changes on different tables or tables without a primary key may become inconsistent.</li> </ul>
	<ul> <li>Note</li> <li>If you use this policy, DDL statements are migrated to Partition 0 of the destination topic by default.</li> <li>If a table does not have a primary key or unique key, DTS migrates the data and DDL statements of the table to Partition 0 of the destination topic.</li> </ul>	

### 5.Precheck and troubleshooting 5.1. Source database connectivity

DTS checks whether DTS servers can connect to the source database during the precheck to ensure successful data migration. This topic describes causes of check failure and how to fix the failure. The following shows the potential causes of the failed source database connectivity check.

#### The database account or password is incorrect

Troubleshooting:

Get a device that can connect to the source database. On the device, enter the database account and password that are specified in the data migration task to check whether the account or password is correct.

**?** Note You can also check the account and password on the server where the source database resides.

#### Solution:

Log on to the DTS console, enter the correct account and password, and perform another precheck.

## The IP address used to access the source database is disallowed to access the database

Troubleshooting:

- You can enter the account and password that are specified in the data migration task on the server where the source database resides and connect to the database. If the connection is successful, it indicates that the source database disallows the access from the IP address.
- If the source database is a MySQL database, you can use a MySQL client to connect to the database and run the following command.

SELECT HOST FROM mysql.user WHERE user='username',password='password';

**?** Note Replace the *username* and *password* with the account and password that are specified in the data migration task.

Check whether the authorized IP address list includes the IP addresses of the DTS servers. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

- If the source database is an SQL Server database, check whether a firewall is set up for the server where the database resides. In addition, check whether the endpoint or triggers in the source database disallows the access from the IP address.
- If the source database is an Oracle database, check whether the TCP.VALIDNODE\_CHECKING it em in the *sqlnet.ora* configuration file is set to yes. If the item is set to yes, it means that the source database disallows the access from the IP address.

Solution:

• If the source database is a MySQL database, run the following command to reauthorize the database account.

GRANT ALL ON . TO 'username'@''%' IDENTIFIED BY 'password';

**?** Note Replace the *username* and *password* with the account and password that are specified in the data migration task.

- If the source database is an SQL Server database, disable the firewall or triggers.
- If the source database is an Oracle database, set the TCP.VALIDNODE\_CHECKING item to no and restart the process.

Log on to the DTS console to perform another precheck.

### A firewall is configured for the server where the source database resides

Troubleshooting:

- If the server where the source database resides runs Windows, find Windows Defender Firewall from the Control Panel to check whether a firewall is configured for the server.
- If the server where the source database resides runs Linux, run the iptables -L command in the shell to check whether a firewall is configured for the server.

Solution:

After you disable the firewall, log on to the DTS console to perform another precheck.

#### Network connections fail

If the connectivity item still cannot pass the precheck after the preceding troubleshooting, network connections between the DTS server and the source database may not function as expected. Submit a ticket to contact Alibaba Cloud engineers.

## 5.2. Destination database connectivity

To ensure successful data migration, DTS checks whether DTS servers can connect to the destination database during precheck. This topic describes the causes of check failures and how to troubleshoot the failures.

A data migration task may fail to pass the connectivity check due to the following reasons.

#### The database account or password is invalid

Troubleshooting:

You can check whether the database account and password for the data migration task are valid from a remote host. However, you must make sure that the remote host can establish a connection to the destination database.

Onte You can also perform the check on the server where the destination database resides.

Solution:

Log on to the DTS console, enter a valid database account and password, and then perform a precheck again.

## The destination database disallows access from external IP addresses

Troubleshooting:

- You can enter the database account and password specified for the data migration task on the server where the destination database resides. This allows you to check whether the server can connect to the destination database. If the server can connect to the destination database, the destination database may disallow access from external IP addresses.
- If the destination database is a MySQL database, you can use a MySQL client to connect to the database and run the following command:

SELECT HOST FROM mysql.user WHERE user='username',password='password';

**?** Note Replace the *username* and *password* with the database account and password that are specified for the data migration task.

Check whether the authorized IP address list includes the CIDR blocks of DTS servers. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

- If the destination database is an SQL Server database, check whether a firewall is configured for the server where the database resides. You can also check whether endpoints or triggers in the destination database disallow access from external IP addresses.
- If the destination database is an Oracle database, check whether the TCP.VALIDNODE\_CHECKING item in the sqlnet.ora configuration file is set to yes. If the item is set to yes, the destination database disallows access from external IP addresses.

Solution:

• If the destination database is a MySQL database, run the following command to authorize the database account again:

GRANT ALL ON \*. \* TO 'username'@''%' IDENTIFIED BY 'password';

(?) Note Replace the *username* and *password* with the database account and password that are specified for the data migration task.

- If the destination database is an SQL Server database, disable the firewall or triggers.
- If the destination database is an Oracle database, set the TCP.VALIDNODE\_CHECKING item to no and restart the process.

Log on to the DTS console to perform a precheck again.

## A firewall is configured for the server where the destination database resides

Troubleshooting:

- If the server where the destination database resides runs on Windows, find Windows Defender Firewall from the Control Panel to check whether a firewall is configured for the server.
- If the server where the destination database resides runs on Linux, run the *iptables -L* command in the shell to check whether a firewall is configured for the server.

Solution:
Disable the firewall, and then log on to the DTS console to perform a precheck again.

### The network is unavailable

If the task still fails to pass the connectivity check after the troubleshooting, the network between DTS servers and the destination database may be unavailable. You can contact technical support by submitting a ticket.

# 5.3. Source database version

To ensure successful data migration, DTS checks the version number of the source database during precheck.

DTS checks whether the source database version meets the requirements. For more information, see Overview of data migration scenarios.

# 6.Migrate data from a selfmanaged database to Alibaba Cloud

# 6.1. Source database: MySQL

# 6.1.1. Migrate data from a self-managed MySQL

# database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a self-managed MySQL database to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a self-managed MySQL database to Alibaba Cloud, you can use all of the supported migration types to ensure service continuity.

# Prerequisites

- An ApsaraDB RDS for MySQL instance is created. For more information, see Create an ApsaraDB RDS for MySQL instance.
- The version of the self-managed MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- The available storage space of the destination ApsaraDB RDS for MySQL instance is larger than the total size of the data in the self-managed MySQL database.

# Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

**?** Note For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database on an ApsaraDB RDS for MySQL instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

# Billing

Migration type	Task configuration feeInternet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# Migration types

#### • Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function.

#### ⑦ Note

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- DTS does not migrate user information. To call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

#### • Full data migration

DTS migrates historical data of the required objects from the self-managed MySQL database to the destination database in the ApsaraDB RDS for MySQL instance.

**?** Note During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.

#### • Incremental data migration

After full data migration is completed, DTS retrieves binary log files from the self-managed MySQL database. Then, DTS synchronizes incremental data from the self-managed MySQL database to the destination ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data from a self-managed MySQL database to Alibaba Cloud.

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statement
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

# Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Self-managed MySQL database	The SELECT permission	The SELECT permission	The REPLICATION SLAVE, REPLICATION CLIENT, SHOW VIEW, and SELECT permissions
ApsaraDB RDS for MySQL instance	The read and write permissions	The read and write permissions	The read and write permissions

For information about how to create and authorize a database account, see the following topics:

- Self-managed MySQL database: Create an account for a user-created MySQL database and configure binary logging
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

# Before you begin

Create an account for a user-created MySQL database and configure binary logging

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

Cancel Set Whitelist and Next

1.Configure Source and Destination	n 2.Configure Migration Types and Objects >	3.Map name modification A.Precheck
* Task Name: M	ysqL_TO_MysqL	]
Source Database		
* Instance Type:	User-Created Database with Public IP Address	v
* Instance Region:	Singapore	Get IP Address Segment of DTS
* Database Type:	MySQL	v
* Hostname or IP Address:		
* Port Number:	3306	
* Database Account:	dtstest	
* Database Password:		${f \Phi}$ Test Connectivity ${igodot}$ Passed
Destination Database		
* Instance Type:	RDS Instance	v
* Instance Region:	Singapore	v
* RDS Instance ID:		•
* Database Account:	dtstest	
* Database Password:		

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to specify a unique task name.
	Instance Type	Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database with Public IP Address</b> .
		<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to configure the <b>Instance Region</b> parameter.
	lnstance Region	<b>Note</b> If a whitelist is configured for the self-managed MySQL database, you must manually add the CIDR blocks of DTS servers to the whitelist of the database. You can click <b>Get IP Address Segment of DTS</b> next to <b>Instance Region</b> to obtain the CIDR blocks of DTS servers.
	Database Type	Select MySQL.
Source		

> Document Version: 20220712

#### Dat a Transmission Service

	Database Section	Parameter	Description
		Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed MySQL database. In this example, enter the public IP address.
		Port Number	Enter the service port number of the self-managed MySQL database. The port must be accessible over the Internet. The default port number is <b>3306</b> .
		Database Account	Enter the account of the self-managed MySQL database. For information about the permissions that are required for the account, see Permissions required for database accounts.
			Enter the password of the database account.
		Database Password	<b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
		Instance Type	Select RDS Instance.
		lnstance Region	Select the region where the destination RDS instance resides.
		RDS Instance ID	Select the ID of the destination RDS instance.
		Dat abase Account	Enter the database account of the destination ApsaraDB RDS for MySQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
			Enter the password of the database account.
	Destinatio n Database	Dat abase Password	<b>?</b> Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
		Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption for an ApsaraDB RDS for MySQL instance.
		~1	<b>Note</b> The Encryption parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### Q Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the objects to be migrated and the migration types.

1.Configu	re Source and	2.Configure Migration T	/pes and	3.Advanced Settings	>	4.Precheck
<ul> <li>Migration T triggers. For m</li> <li>Note: do no cleans up th</li> <li>Data migrat</li> </ul>	ypes: Schema Migratio ore information, see Refer t clean up the incremental e log too early, the DTS in ion applies to short-term n	on Full Data Migration ence data log generated by the sour cremental task may fail igration scenarios. Typical scer	Incremental E ce database after the arios include migratin	Data Migration Note: Incre DTS task is started when the D	mental data migra )TS full task is runi d sharding databas	tion does not support ning. If the source database ses, and migrating data
between Ap For long-ter Available	sara Stack databases, m data synchronization in i	real time, use the data synchron	nization feature.	Selected (To edit an object	name or its filter,	hover over the object and click
Expand the	tree before you perform a estdata ables iews	i gloi   Q	> <	dtstestdata (201	bjects)	
Select All				Remove All		
*Rename Data	bases and Tables:	Do Not Change Databas	e and Table Names	<ul> <li>Change Database and T</li> </ul>	îable Names	
* Retry Time f *Source table want to copy t	or Failed Connection DMSONLINEDo you he temporary table to	720 Minute	5 🕜			
Information: 1. Data migrat in the source d 2. Do not do D	on only copies the data an atabase. DL operation during struct	d schema in the source databas ure and full migration, otherwis	e and saves the copy e the task may fail	in the destination database. T	he process does no	ot affect any data or schema
				Cancel	Previous	Save Precheck
Setting	Description					
	<ul> <li>To perform</li> <li>Migration</li> </ul>	n only full data mi	gration, sele	ct Schema Migra	tion and F	ull Data
Select the	<ul> <li>To ensure</li> <li>Data Migr</li> </ul>	service continuity or a continuity of a contin	during data r nental Data	nigration, select S Migration.	chema Mi	gration, Full
nigratio						

**Note** If **Incremental Data Migration** is not selected, do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.

n types

# Data Migration Migrate data from a self-managed database to Alibaba Cloud

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the > icon to add the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated. If you select tables or columns as the objects to be migrated, DTS does not migrate other objects such as views, triggers, and stored procedures to the destination database.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination RDS instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS reconnects to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
е	

Setting	Description
Specify whether to copy tempora	If you use <mark>Data Management (DMS)</mark> to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations.
ry tables	• Yes: DTS migrates the data of temporary tables generated by online DDL operations.
to the destinati on databas e when DMS perform s online DDL operatio ns on the source table	<b>Note</b> If online DDL operations generate a large amount of data, the migration task may be delayed.
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
	ONDE If you select No, the tables in the destination database may be locked.

#### 8. Click Precheck.

? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.

#### Stop the migration task

**Warning** We recommend that you prepare a rollback solution to migrate incremental data from the destination database to the source database in real time. This allows you to minimize the negative impact of switching your workloads to the destination database. For more information, see Switch workloads to the destination database. If you do not need to switch your workloads, you can perform the following steps to stop the migration task.

• Full data migration

Do not manually stop a task during full data migration. Otherwise, the system may fail to migrate all data. Wait until the migration task automatically ends.

Incremental data migration

The task does not automatically end during incremental data migration. You must manually stop the migration task.

- i. Wait until the task progress bar shows Incremental Data Migration and The migration task is not delayed. Then, stop writing data to the source database for a few minutes. In some cases, the progress bar shows the delay time of incremental data migration.
- ii. After the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.

Task Name	✓ Search by migration task name.	Search Sort: Default Sorting	✓ Status: All	✓ 🐌 Tag	
Task ID/I	Name:	Status: Migrating Quick Diagnostics   F	Pause Task   View Details   Dup	uplicate Task Upgrade   Configure Monito	ring and Alerting   Change password   Edit Tag
4 Nov 20 Schem	021, 09:59:01 Created na Migration <b>100%</b>	Full Data Migration 100%(Migra	nted Rows: 1)	Incremental Data M	Completed
☑ 3	Pause Stop Delete Edit Tag			Total: 1 item(s),	Per Page: 20 item(s)

### What to do next

The database accounts that are used for data migration have the read and write permissions. After you migrate data, you must delete the database accounts to ensure security.

# FAQ

- Q: What can I do if a migration task fails to pass the precheck? A: For more information, see Source database connectivity.
- Q: How can I troubleshoot a failed migration task? A: For more information, see Fix a failed data migration task.

# 6.1.2. Migrate data from a self-managed MySQL database connected over Express Connect, VPN

# Gateway, or Smart Access Gateway to an

# ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a self-managed MySQL database that is connected over Express Connect, VPN Gateway, or Smart Access Gateway to an ApsaraDB RDS for MySQL database by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a self-managed MySQL database, you can select all of the supported migration types to ensure service continuity.

### Prerequisites

- The version of the self-managed MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total size of the data in the self-managed MySQL database.
- The on-premises network to which the self-managed MySQL database belongs is connected to Alibaba Cloud over Express Connect, VPN Gateway, or Smart Access Gateway.

⑦ Note For more information, see Connect an on-premises database to DTS by using CEN.

### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

⑦ Note For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database on an ApsaraDB RDS for MySQL instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

### Billing

Migration type	Task configuration fee Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# **Migration types**

• Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function.

? Note

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- DTS does not migrate user information. To call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.
- Full data migration

DTS migrates historical data of the required objects from the self-managed MySQL database to the destination database in the ApsaraDB RDS for MySQL instance.

**?** Note During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.

#### • Incremental data migration

After full data migration is completed, DTS retrieves binary log files from the self-managed MySQL database. Then, DTS synchronizes incremental data from the self-managed MySQL database to the destination ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data from a self-managed MySQL database to Alibaba Cloud.

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statement
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

#### Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Self-managed MySQL database	The SELECT permission	The SELECT permission	The REPLICATION SLAVE, REPLICATION CLIENT, SHOW VIEW, and SELECT permissions
ApsaraDB RDS for MySQL instance	The read and write permissions	The read and write permissions	The read and write permissions

For more information about how to create and authorize a database account, see the following topics:

- Self-managed MySQL database: Create an account for a user-created MySQL database and configure binary logging
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

#### Before you begin

- 1. Create an account for a user-created MySQL database and configure binary logging.
- 2. Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.

1.Configure Source and Destination D	Databases 2. Configure Migration Types and Objects	3.Advanced Settings	A.Precheck
Task Name:	MySQL_TO_RDS MySQL		
Source Database			
* Instance Type:	User-Created Database Connected Over Express Connect, VPP 🔻	DTS support type	
* Instance Region:	China (Hangzhou)	Guide	
* Peer VPC:	vpc-bp	Proprietary network of Other Apsara Stack Accounts	
* Database Type:	MySQL		
* IP Address:	172.16.88.	]	
* Port Number:	3306		
* Database Account:	dtstest		
* Database Password:		Test Connectivity	
Destination Database			
* Instance Type:	RDS Instance v		
* Instance Region:	China (Hangzhou)		
* RDS Instance ID:	rm-1u v		
* Database Account:	dtstest		
Database Password:	······	Test Connectivity	
* Encryption:	Non-encrypted     SSL-encrypted		
			Cancel Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select User-Created Database Connected over Express Connect, VPN Gateway, or Smart Access Gateway.
	Instance Region	Select the region to which the virtual private cloud (VPC) that is connected to Express Connect, VPN Gateway, or Smart Access Gateway belongs.
	Peer VPC	Select the VPC that is connected to Express Connect, VPN Gateway, or Smart Access Gateway.
	Database Type	Select MySQL.
	IP Address	Enter the endpoint that is used to access the self-managed MySQL database.

Section Parameter Source		Description		
Database	Port Number	Enter the service port number of the self-managed MySQL database. The default port number is <b>3306</b> .		
	Database Account	Enter the account of the self-managed MySQL database. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
		Enter the password of the database account.		
	Dat abase Password	<b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		
	Instance Type	Select RDS Instance.		
	Instance Region	Select the region where the destination RDS instance resides.		
	RDS Instance ID	Select the ID of the destination RDS instance.		
	Dat abase Account	Enter the database account of the destination RDS instance. For m information about the permissions that are required for the accou see Permissions required for database accounts.		
		Enter the password of the database account.		
Destinatio n Database	Dat abase Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		
		Coloct Non ancrupted or CCL ancrupted If you want to coloct		
	Encryption	SSL-encrypted of SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption on an ApsaraDB RDS for MySQL instance.		
		regions in mainland China and the China (Hong Kong) region.		

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

1.Configu	ire Source and 🔶 2	2.Configure Migration Typ	es and	3.Advanced Sett	ings >	4.Precheck
<ul> <li>Migration 1 triggers. For m Note: do no cleans up th Data migrat between Ap For long-ter</li> </ul>	ypes: ♥ Schema Migration ore information, see Reference t clean up the incremental data e log too early, the DTS increm ion applies to short-term migra sara Stack databases. m data synchronization in real	Full Data Migration log generated by the source rental task may fail tion scenarios. Typical scenar time, use the data synchroniz	✓ Incremental C database after the ios include migratin ation feature.	iata Migration Not	e: Incremental data mign en the DTS full task is rur aling and sharding databa	ation does not support ming. If the source database uses, and migrating data
Available Expand the e e dtst e e T e e V	: tree before you perform a glo estdata ables 'iews		> <	Selected (To edit a Edit.) Learn more.	n object name or its filter a (20bjects)	hover over the object and click
Select All				Remove All		
*Rename Data * Retry Time f *Source table want to copy to the target data Information: 1. Data migrat in the source of 2. Do not do D	bases and Tables: () or Failed Connection DMS_ONLINE_Do you () he temporary table to base during DDL: ion only copies the data and so latabase. DL operation during structure a	Do Not Change Database a     720     Minutes     Yes     No     O	and Table Names          Image: same same same same same same same same	Change Databases of the contract of the destination datases of the desti	se and Table Names base. The process does r	iot affect any data or schema
				Car	Previous	Save Precheck
Setting	Description					
<ul> <li>To perform only full migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full</li> <li>Data Migration, and Incremental Data Migration.</li> </ul>						

**Notice** If **Incremental Data Migration** is not selected, we recommend that you do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.

migratio

n types

#### Dat a Transmission Service

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Notice</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
the source or destinati on databas	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
e	

Setting	Description
Specify whether to copy tempora ry tables to the doctinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.
on databas e when DMS	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.
perform s online ומס	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
operatio ns on the source	<b>Note</b> If you select No, the tables in the destination database may be locked.

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of incremental data migration changes to The migration task is not delayed again. Then, manually stop the migration task.



12. Switch your workloads to the ApsaraDB RDS for MySQL instance.

# 6.1.3. Migrate data from a self-managed MySQL database connected over Express Connect, VPN Gateway, or Smart Access Gateway to an ApsaraDB RDS for MySQL instance across Alibaba

# **Cloud accounts**

This topic describes how to migrate data from a self-managed MySQL database that is connected over Express Connect, VPN Gateway, or Smart Access Gateway to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). In this scenario, the Express Connect circuit and the destination RDS instance are owned by different Alibaba Cloud accounts. DTS supports schema migration, full data migration, and incremental data migration. You can select all of the supported migration types to ensure service continuity.

#### Prerequisites

- The version of the self-managed MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total size of the data in the self-managed MySQL database.
- The on-premises network to which the self-managed MySQL database belongs is connected to Alibaba Cloud VPC over Express Connect, VPN Gateway, or Smart Access Gateway. The Express Connect circuit and the destination RDS instance are owned by different Alibaba Cloud accounts.

⑦ Note For more information, see Connect an on-premises database to DTS by using CEN.

#### Context

The data center that hosts your database is connected to Alibaba Cloud VPC over Express Connect, VPN Gateway, or Smart Access Gateway. You need to migrate data from the on-premises database to

an ApsaraDB RDS for MySQL instance across different Alibaba Cloud accounts. The following figure shows the architecture for this scenario.



#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

(?) Note For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database on an ApsaraDB RDS for MySQL instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source

instance will overwrite the data in the destination instance after the task is resumed.

# Billing

Migration type Task configuration fee		Internet traffic fee		
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from		
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.		

# **Migration types**

#### • Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function.

#### ⑦ Note

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- DTS does not migrate user information. To call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

#### • Full data migration

DTS migrates historical data of the required objects from the self-managed MySQL database to the destination database in the ApsaraDB RDS for MySQL instance.

**?** Note During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is completed, the tablespace of the destination database is larger than that of the source database.

#### • Incremental data migration

After full data migration is completed, DTS retrieves binary log files from the self-managed MySQL database. Then, DTS synchronizes incremental data from the self-managed MySQL database to the destination ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data from a self-managed MySQL database to Alibaba Cloud.

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statement
DML	INSERT, UPDATE, DELETE, and REPLACE

Operatio n type	SQL statement
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

### Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Self-managed MySQL database	The SELECT permission	The SELECT permission	The REPLICATION SLAVE, REPLICATION CLIENT, SHOW VIEW, and SELECT permissions
ApsaraDB RDS for MySQL instance	The read and write permissions	The read and write permissions	The read and write permissions

For information about how to create and authorize a database account, see the following topics:

- Self-managed MySQL database: Create an account for a user-created MySQL database and configure binary logging
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

### Before you begin

- 1. Create an account for a user-created MySQL database and configure binary logging.
- 2. Log on to the Alibaba Cloud Management Console by using the Alibaba Cloud account that owns the Express Connect circuit. Authorize DTS to access the network that is connected over Express Connect. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.
- 3. Create a RAM role and authorize the RAM role to access the resources of the Alibaba Cloud account. For more information, see Configure RAM authorization for data migration or synchronization from a self-managed database in a VPC across different Alibaba Cloud accounts.

#### Procedure

- 1. Use the Alibaba Cloud account that owns the destination RDS instance to log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Select User-Created Database Connected over Express Connect, VPN Gateway, or Smart Access Gateway as the instance type. Then, click VPC of Another Alibaba Cloud Account next to the Peer VPC field.

#### Dat a Transmission Service

1.Configure Source and Destination	2.Configure Migration Types and Objects	3.	Advanced Settings	>	4.Precheck
* Task Name: 👔	M sh Ge				
Source Database					
* Instance Type:	User-Created Database with Public IP Address	۳	DTS support type		
* Instance Region:	China (Hangzhou)	•	Get IP Address Segment of	DTS	
* Database Type:	MySQL	•			
* Hostname or IP Address:					
* Port Number:	3306				
* Database Account:					
* Database Password:			Test Connectivity		

6. Configure the source and destination databases.

<b>*</b> Ta	sk Name: 👔	ſγSQL			
Source Database	Source Database				
* Ins	Instance Type: User-Created Database Connected Over Express Connect, VP↑		DTS support type		
* Insta	ince Region:	China (Hangzhou)	Ŧ	Guide	
*Apsara Stack Tenant	Account ID:				
	Role Name:	ram-for-dts		Authorize Role Across Accounts	
	* Peer VPC:		-	Proprietary network of the current login account	
* Dat	abase Type:	MySQL	Ŧ		
*	IP Address:	172.16.		]	
<b>*</b> p	ort Number:	3306			
* Databa	ase Account:	dtstest		]	
* Databas	e Password:	•••••	<i><b>(</b></i> )	Test Connectivity	
Destination Database	Destination Database				
* Ins	stance Type:	RDS Instance	Ŧ		
* Instance Region:		China (Hangzhou)			
* RDS Instance ID: n		rm-	•		
* Databa	ase Account:	dtstest		]	
* Databas	e Password:	•••••	4»	Test Connectivity	
*	Encryption:	Non-encrypted O SS	L-encrypted		
					Cancel Set Whitelist and Next
Section	Para	meter	Description		
N/A	Tasl	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.		mmend that you You do not need	
	Insta	ance Type	Select User-Create Connect, VPN Gat	d Database Connected over eway, or Smart Access Gate	Express way.

Section	Parameter	Description			
	Instance Region	Select the region of the VPC that is connected to the self-managed MySQL database.			
		Enter the ID of the Alibaba Cloud account that owns the Express Connect circuit.			
	Alibaba Cloud Account ID	<b>Note</b> To obtain the ID of the Alibaba Cloud account that owns the Express Connect circuit, you must log on to the Account Management console by using this account. The account ID is displayed on the Security Settings page.			
		Account Management         Account Center > Basic Information         Copin Account         ort**@last aligned cent         Login Account         ort**@last aligned cent         Vertified         No real-name submittation         Copin Account         Account ID         Login Account ID         Control To         Contro         Control To         Control			
Source Database	Role Name	Enter the name of the RAM role that you created earlier in Before you begin.			
	Peer VPC	Select the ID of the VPC that is connected to the self-managed MySQL database.			
	Database Type	Select MySQL.			
	IP Address	Enter the endpoint that is used to access the self-managed MySQL database.			
	Port Number	Enter the service port number of the self-managed MySQL database. The default port number is <b>3306</b> .			
	Database Account	Enter the account of the self-managed MySQL database. For more information about the permissions that are required for the account, see Permissions required for database accounts.			
		Enter the password of the database account.			
	Database Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			
	Instance Type	Select RDS Instance.			
	Instance Region	Select the region where the destination RDS instance resides.			
	RDS Instance ID	Select the ID of the destination RDS instance.			

Section	Parameter	Description	
	Dat abase Account	Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.	
		Enter the password of the database account.	
Destinatio n Database	Dat abase Password	<b>?</b> Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.	
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see <b>Configure SSL encryption on an ApsaraDB RDS for MySQL instance</b> .	
	спетурноп	<b>Note</b> The Encryption parameter is available only for regions in mainland China and the China (Hong Kong) region.	

#### 7. In the lower-right corner of the page, click Set Whitelist and Next.

Once DTS adds the CIDR blocks of DTS servers to the whitelist of the destination ApsaraDB RDS for MySQL instance. This ensures that DTS servers can connect to the destination RDS instance.

#### 8. Select the migration types and the objects to be migrated.

1.comigu	re Source and	2.Configure Migration Typ	pes and 🔪	3.Advanced Settings	4.Precheck
<ul> <li>Migration Tr triggers. For more Note: do not cleans up the Data migrati between Aps For long-terr</li> </ul>	ypes: Schema Mig ore information, see Re t dean up the incremer e log too early, the DT ion applies to short-ten sara Stack databases, m data synchronization	ration Full Data Migration sference tal data log generated by the source S incremental task may fail m migration scenarios. Typical scena in real time, use the data synchroni	✓ Incremental C a database after the rios include migratin zation feature.	Data Migration Note: Incremental data mi DTS task is started when the DTS full task is r Ing data to the doud, scaling and sharding data	gration does not support unning. If the source database bases, and migrating data
Available Expand the end of the formation of the formation end of the formation of the form	tree before you perfor estdata ables iews	m a gloi I Q	> <	Selected (To edit an object name or its filt Edit.) Learn more.	er, hover over the object and click
Select All				Remove All	
*Rename Datal * Retry Time for *Source table I want to copy th the target datal Information: 1. Data migrati in the source d 2. Do not do DI	bases and Tables: or Failed Connection DMS_ ONLINE_ Do you te temporary table to base during DDL: on only copies the data atabase. DL operation during str	Do Not Change Database     720     Minutes     Yes     No     O	and Table Names       Image: mail of the second	Change Database and Table Names	not affect any data or schema
				Cancel Previous	Save Precheck

• To ensure service continuity during data migration, select **Schema Migration**, **Full Data Migration**, and **Incremental Data Migration**.

**?** Note If Incremental Data Migration is not selected, we recommend that you do not write data to the source database during data migration. This ensures data consistency between the source and destination databases.

Select

n types

the migratio

#### Dat a Transmission Service

Setting	Description
Select the objects to be migrate d	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
e	

Setting	Description			
Specify whether to copy tempora ry tables to the destinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.			
on databas e when DMS perform s online DDL operatio ns on the	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.			
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.			
	<b>Note</b> If you select No, the tables in the destination database may be locked.			
table				

#### 9. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.

#### 10. After the task passes the precheck, click Next.

- 11. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 12. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of incremental data migration changes to The migration task is not delayed again. Then, manually stop the migration task.

Task Name   Search by migration task name.	Search Sort: Default Sorting v Status: All v	✤ Tag
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate	ask Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag Completed
Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Incremental Data M in The migration task is not delayed. Total: 1 item(s), Per Page: 20 item(s) <

13. Switch your workloads to the destination ApsaraDB RDS for MySQL instance.

# 6.1.4. Migrate data from a self-managed MySQL database to a PolarDB for MySQL cluster

PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. PolarDB is compatible with the MySQL database engine and features high availability, ease of use, and reliability. This topic describes how to migrate data from a self-managed MySQL database to a PolarDB for MySQL cluster by using Data Transmission Service (DTS).

### Prerequisites

- The engine version of the self-managed MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- A PolarDB for MySQL cluster is created. For more information, see Create a PolarDB for MySQL cluster.
- If the source MySQL database is an on-premises database, the CIDR blocks of DTS servers must be added to the IP whitelist of the database. This setting ensures that DTS servers can access the source MySQL database. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or

DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.

• If a data migration task fails, DTS automatically resumes the task. Therefore, before you switch your workloads to the destination cluster, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination cluster after the task is resumed.

## Migration types

DTS supports schema migration, full data migration, and incremental data migration. For more information, see Terms.

**Note** When you migrate data between PolarDB for MySQL clusters, you can select all of the supported migration types to ensure service continuity.

# Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.	

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

### Permissions required for database accounts

Database	Schema migration and full data migration	Incremental data migration
Self-managed MySQL database	The SELECT permission	The REPLICATION SLAVE, REPLICATION CLIENT, SHOW VIEW, and SELECT permissions
PolarDB for MySQL cluster	Read and write permissions	Read and write permissions

For more information about how to create a database account and grant permissions to the account, see the following topics:

- Self-managed MySQL databases: Create an account for a user-created MySQL database and configure binary logging
- PolarDB for MySQL clusters: Create a database account

#### Preparations

Create an account for a user-created MySQL database and configure binary logging

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination 2.Configure Migration Types and Objects 3.Map name modification 4.Precheck					
Task Name: MySQL_TO_POLARD8					
Source Database					
* Instance	Type: User-Created Database wit	h Public IP Address	v		
* Instance Re	gion: Singapore		Get IP Address Segment of DTS		
* Database 1	Type: MySQL		Y		
* Hostname or IP Add	dress:				
* Port Nur	mber: 3306				
* Database Acc	ount: dtstest				
* Database Passi	word:	đ.			
Destination Database					
* Instance T	Type: PolarDB		T		
* Instance Re	gion: Singapore	,	T		
* PolarDB Instanc	e ID:		•		
* Database Acc	ount: dtstest				
* Database Passi	word:	đ.			
	Cancel Set Whitelist and Next				
Section	Parameter	Description			
None Task Name The task name that that you specify a d the task. You do no		DTS automatically generates. We re descriptive name that makes it easy t t need to specify a unique task nam	ecommend o identify e.		

Section	Parameter	Description		
	Instance Type	The instance type of the source database. In this example, <b>User-</b> <b>Created Database with Public IP Address</b> is selected for this parameter.		
		<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .		
		If you select <b>User-Created Database with Public IP Address</b> for the instance type, you do not need to configure the <b>Instance Region</b> parameter.		
	Instance Region	<b>? Note</b> If a whitelist is configured for the self-managed MySQL database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click <b>Get IP Address Segment of DTS</b> next to <b>Instance Region</b> to obtain the CIDR blocks of DTS servers.		
	Database Type	Select MySQL.		
	Hostname or IP Address	The endpoint that is used to connect to the self-managed MySQL database. In this example, the public IP address is used.		
Source Database	Port Number	The service port number of the self-managed MySQL database. Default value: <b>3306</b> .		
	Dat abase Account	The account of the self-managed MySQL database. For information about the permissions that are required for the account, see Permissions required for database accounts.		
		The password of the database account.		
	Dat abase Password	Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message is displayed. If the Failed message is displayed, click Check next to Failed. Modify the source database parameters based on the check results.		

Section	Parameter	Description				
Destination Instance Details	Instance Type	Select <b>PolarDB</b> .				
	Instance Region	The region where the destination PolarDB cluster resides.				
	PolarDB Instance ID	The ID of the destination PolarDB for MySQL cluster.				
	Dat abase Account	The database account of the destination PolarDB cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.				
	Dat abase Password	The password of the database account.				
		<b>Note</b> After you specify the destination database parameters, click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message is displayed. If the <b>Failed</b> message is displayed, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.				

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### ☐ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.

7. Select the migration types and the objects to be migrated.

1.Configu	ire Source and 2.0	Configure Migration Type	es and	3.Advanced	Settings	>	4.Prechec	¢
<ul> <li>Migration T triggers. For m</li> <li>Note: do not cleans up th</li> <li>Data migrat between Ap For long-ten</li> </ul>	ypes: Schema Migration ore information, see Reference t clean up the incremental data k le log too early, the DTS increment ion applies to short-term migration sara Stack databases. m data synchronization in real tin	Full Data Migration g generated by the source- ital task may fail n scenarios. Typical scenari re, use the data synchronizz	Incremental D database after the l os include migrating ation feature.	ata Migration DTS task is starts g data to the clou	Note: Increment and when the DTS and, scaling and sh	ntal data migrat full task is runn narding database	ion does not supp ing. If the source es, and migrating	iort database data
Available Expand the tree before you perform a glo Q  C C C C C C C C C C C C C C C C C C		> <	Selected (To edit an object name or its filter, hover over the object an Edit.) Learn more. Q  d dtstestdata (20bjects) customer order			ect and click		
Select All *Rename Data * Retry Time fr *Source table	bases and Tables: or Failed Connection DMS_ONLINE_Do you	Do Not Change Database a 720 Minutes Yes () No (?)	nd Table Names	Remove All Change Da	atabase and Tabl	e Names		
want to copy the temporary table to the target database during DDL: Information: 1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database. 2. Do not do DDL operation during structure and full migration, otherwise the task may fail								
Setting	Description				Cancel	Previous	Jave	
Select	<ul> <li>To perform only full migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> </ul>							

**Notice** If **Incremental Data Migration** is not selected, we recommend that you do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.

the

migratio

n types

#### Dat a Transmission Service

Setting	Description						
Select the objects to be migrate d	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.						
	<ul> <li>Notice</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>						
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.						
Specify the retry time for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.						
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.						
Setting	Description						
--	---	--	--	--	--		
Specify whether to copy tempora ry tables to the doctinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.						
on databas e when DMS	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.						
perform s online	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.						
operatio ns on the source table	<b>Note</b> If you select No, the tables in the destination database may be locked.						
taple							

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.

Task Name  V Search by migration task name.	Search Sort: Default Sorting V Status: All V	Tag
Task ID/Name:     A Nov 2021, 09:59:01 Created	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag Completed
Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Incremental Data M 1 n The migration task is not delayed. Total: 1 item(s), Per Page: 20 item(s)

12. Switch your workloads to the destination PolarDB cluster.

# 6.1.5. Migrate data from a self-managed MySQL database to a PolarDB-X instance

This topic describes how to migrate data from a self-managed MySQL database to a PolarDB-X 1.0 instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. You can select all of the supported migration types to ensure service continuity when you migrate data from the self-managed Oracle database to the PolarDB-X 1.0 instance.

## Prerequisites

- The version of the self-managed MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- The databases in the PolarDB-X 1.0 instance are created based on ApsaraDB RDS for MySQL instances. DTS does not support DRDS databases that are created based on PolarDB for MySQL clusters.
- The available storage space of the RDS instances in the PolarDB-X 1.0 instance is larger than the total size of the data in the self-managed MySQL database.

## Precautions

• DTS does not support schema migration from a self-managed MySQL database to a PolarDB-X 1.0 instance.

**?** Note During schema migration, DTS migrates the schemas of the required objects, such as tables, from the source database to the destination database.

 DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.

- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration Free of charge.		Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	Alibaba Cloud over the Internet. For more information, see Pricing.

## **Migration types**

• Full data migration

DTS migrates historical data of the required objects from the self-managed MySQL database to the destination database in the PolarDB-X 1.0 instance.

**?** Note During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination instance. After full data migration is complete, the tablespace of the destination instance is larger than that of the source database.

#### • Increment al data migration

After full data migration is complete, DTS retrieves binary log files from the self-managed MySQL database. Then, DTS synchronizes incremental data from the self-managed MySQL database to the destination PolarDB-X 1.0 instance. Incremental data migration allows you to ensure service continuity when you migrate data from the self-managed MySQL database to the PolarDB-X 1.0 instance.

## SQL operations that can be synchronized during incremental data migration

INSERT, UPDATE, DELETE, and REPLACE operations

## Permissions required for database accounts

Database	Full data migration	Incremental data migration
Self-managed MySQL database	The SELECT permission	The REPLICATION CLIENT, REPLICATION SLAVE, SHOW VIEW, and SELECT permissions

Database	Full data migration	Incremental data migration
PolarDB-X 1.0	The read and write permissions	The read and write permissions

## Before you begin

- 1. Create an account for a user-created MySQL database and configure binary logging.
- 2. Create a database and tables in the destination PolarDB-X 1.0 instance based on the schema of the self-managed MySQL database. For more information, see Create a DRDS database and Create a DRDS table.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the **Migration Tasks** page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination Databases 2.Configure Migration Types and Objects		$\geq$	> 3.Map nam	e modification	>		4.Precheck		
* Task	Name: M	lySQL_TO_DRDS							
Source Database									
* Insta	nce Type:	User-Created Datab	ase with Public IP Address	Ŧ					
* Instan	e Region:	Singapore	ale with rabile in Address	*	Get IP Address Seament	of DTS			
* Datab	ase Type:	Singapore ·		Ŧ					
* Hostname or If	P Address:								
* Por	t Number:	3306							
* Database	Account:	dtstest							
* Database	Database Password:		<	<b>\$</b> >	Test Connectivity	⊘ Passed			
Destination Database									
* Insta	nce Type:	DRDS Instance		Ŧ					
* Instanc	e Region:	Singapore		Ŧ					
* DRDS In:	stance ID:			Ŧ					
* Databa	ise Name:	mysqltest 🔻							
* Database	e Account:	dtstest							
* Database	Password:	•••••	<	<b>\$</b> >	Test Connectivity	⊘ Passed			
								Cancel	Set Whitelist and Next
Section	Par	ameter	Description						
N/A	Tas	ik Name	DTS automatical specify an inform to use a unique t	lly na tas	generates a t tive name for sk name.	ask name. easy iden	We recontification.	mmenc You de	l that you o not need

Section	Parameter	Description			
	Instance Type	Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database with Public IP Address</b> .			
		<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .			
		If the instance type is set to <b>User-Created Database with Public</b> <b>IP Address</b> , you do not need to specify the <b>instance region</b> .			
	Instance Region	Note If a whitelist is configured for the self-managed MySQL database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.			
	Database Type	Select MySQL.			
Source Database	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed MySQL database. In this example, enter the public IP address.			
	Port Number	Enter the service port number of the self-managed MySQL database. The default port number is <b>3306</b> .			
	Dat abase Account	Enter the account of the self-managed MySQL database. For more information about the permissions that are required for the account, see Permissions required for database accounts.			
	Dat abase Password	Enter the password of the database account.			
		<b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			
	Instance Type	Select DRDS Instance.			
	Instance Region	Select the region where the destination PolarDB-X 1.0 instance resides.			

Section	Parameter	Description		
	DRDS Instance ID	Select the ID of the destination PolarDB-X 1.0 instance.		
Destinatio n	Dat abase Name	Select the name of the destination database.		
Database	Dat abase Account	Enter the database account of the destination PolarDB-X 1.0 instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
		Enter the password of the database account.		
	Dat abase Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

Note DTS adds the CIDR blocks of DTS servers to the whitelist of the destination
 PolarDB-X 1.0 instance. This ensures that DTS servers can connect to the destination PolarDB-X
 1.0 instance.

7. Select the migration types and the objects to be migrated.

1.Configu	ure Source and	2.Configure Migration Typ	pes and	3.Advanced Settings	>	4.Precheck
<ul> <li>Migration T triggers. For m</li> <li>Note: do not cleans up th</li> <li>Data migrati</li> <li>between Ap;</li> <li>For long-ten</li> </ul>	Types: Schema Migration nore information, see Reference at clean up the incremental dat the log too early, the DTS increment ion applies to short-term migra isara Stack databases. m data synchronization in real	Full Data Migration a log generated by the source mental task may fail ation scenarios. Typical scena time, use the data synchroni	✓ Incremental D a database after the I rios include migratin zation feature.	ata Migration Note: Increm DTS task is started when the DTS g data to the cloud, scaling and s	ental data migratic S full task is runnir sharding database:	n does not support ng. If the source database s, and migrating data
Available Expand the e dtst	e tree before you perform a glo testdata Tables Views		> <	Selected (To edit an object na Edit.) Learn more.	ects)	over over the object and click
Select All				Remove All		
*Rename Data * Retry Time for *Source table I want to copy th the target data Information: 1. Data migrati in the source d 2. Do not do D	abases and Tables: for Failed Connection DM5_ONLINE_ Do you he temporary table to abase during DDL: ion only copies the data and so database. DDL operation during structure	Do Not Change Database     720     Minutes     Yes     No     O	and Table Names ② and saves the copy the task may fail	Change Database and Tak in the destination database. The	process does not	affect any data or schema
				Cancel	Previous	Save Precheck
etting	Description					
	• To perform o	only full data mig	ration, seled	t only <b>Full Data M</b>	ligration.	

• To ensure service continuity during data migration, select both **Full Data Migration** and **Incremental Data Migration**.

**?** Note If Incremental Data Migration is not selected, do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.

Select the

migratio n types

#### Dat a Transmission Service

Setting	Description
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or databases as the objects to be migrated. If you select tables or columns as the objects to be migrated, DTS does not migrate other objects such as views, triggers, and stored procedures to the destination database.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>
Specify whether to rename object names	You can use the object name mapping feature to change the names of the objects that are synchronized to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

Setting	De	Description					
Specify whether to copy tempora ry tables to the	In Hc us yo o	this scenario, DTS does not migrate DDL operations. You must set this parameter to No. wever, DTS may support DDL operations for this scenario in the future. In this case, if you e Data Management (DMS) to perform online DDL operations on the source database, u can specify whether to migrate temporary tables generated by the operations. Yes: DTS migrates the data of temporary tables generated by online DDL operations.					
destinati on databas e when		<b>Note</b> If online DDL operations generate a large amount of data, the migration task may be delayed.					
DMS perform s online	0	<b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.					
DDL operatio ns on		<b>Note</b> If you select No, the tables in the destination database may be locked.					
the source table							

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

sk Name	Search Sort: Default Sorting V Status: All V 🗣 1	ag
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed
3 Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) <pre>« &lt; 1 &gt; &gt;</pre>

12. Switch your workloads to the destination PolarDB-X 1.0 instance.

## What to do next

The database accounts that are used for data migration have the read and write permissions. After data migration is complete, you must delete the accounts of both the self-managed MySQL database and the PolarDB-X 1.0 instance to ensure security.

## 6.2. Source database: SQL Server

## 6.2.1. Migrate incremental data from a self-

## managed SQL Server database to an ApsaraDB RDS for SQL Server instance

This topic describes how to migrate incremental data from a self-managed SQL Server database to an ApsaraDB RDS for SQL Server instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a self-managed SQL Server database, you can select all of the supported migration types to ensure service continuity.

**Note** For information about how to perform only full data migration, see Migrate full data from a self-managed SQL Server database to an ApsaraDB RDS for SQL Server instance.

## Prerequisites

• The version of the self-managed SQL Server database is 2008, 2008 R2, 2012, 2014, 2016, 2017, or 2019.

#### ? Note

- If a primary/secondary switchover is performed in an SQL Server cluster or Always On availability group (AOAG), the log serial numbers of the primary database and the secondary database become inconsistent. In this case, DTS considers that the logs of the source database are discontinuous, and the migration task fails. Therefore, you cannot use an SQL Server cluster or an SQL Server AOAG as the source database.
- If you migrate data between different versions of databases, make sure that the database versions are compatible.
- The tables to be migrated from the self-managed SQL Server database have primary keys or UNIQUE NOT NULL indexes.
- The available storage space of the ApsaraDB RDS for SQL Server instance is larger than the total size of the data in the self-managed SQL Server database.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- To ensure that the incremental data migration task runs as expected, do not frequently back up the source database. We recommend that you retain log files for more than three days. Otherwise, you cannot retrieve log files after they are truncated.
- To ensure that the delay time of incremental data migration is accurate, DTS adds a heartbeat table to the self-managed SQL Server database. The name of the heartbeat table is Source table name\_dts\_mysql\_heartbeat .
- DTS automatically creates a destination database in the ApsaraDB RDS for SQL Server instance. However, if the name of the source database is invalid, you must create a database in the ApsaraDB RDS for SQL Server instance before you configure the data migration task.

(?) Note For more information about the naming conventions of ApsaraDB RDS for SQL Server databases and how to create a database, see Create a database on an ApsaraDB RDS for SQL Server instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

#### Limits

- DTS does not migrate the schemas of assemblies, service brokers, full-text indexes, full-text catalogs, distributed schemas, distributed functions, CLR stored procedures, CLR scalar-valued functions, CLR table-valued functions, internal tables, systems, or aggregate functions.
- DTS does not migrate data of the SQL\_VARIANT type.

- DTS does not migrate tables that contain computed columns.
- A single data migration task can migrate incremental data from only one database. To migrate incremental data from multiple databases, you must create a data migration task for each database.

## **Migration types**

• Schema migration

DTS migrates the schemas of the required objects to the destination database. DTS supports schema migration for the following types of objects: table, view, trigger, synonym, SQL stored procedure, SQL function, plan guide, user-defined type, rule, default, and sequence.

- Full data migration DTS migrates historical data of the required objects from the self-managed SQL Server database to the destination database.
- Incremental data migration
   After full data migration is complete, DTS migrates incremental data from the self-managed SQL
   Server database to the destination database.

## SQL operations that can be synchronized during incremental data migration

• INSERT, UPDATE, and DELETE

**?** Note If an UPDATE operation updates only the large fields, DTS does not synchronize the operation.

• CREATE TABLE

**?** Note If a CREATE TABLE operation creates a partitioned table or a table that contains functions, DTS does not synchronize the operation.

- ALTER TABLE operations, including only ADD COLUMN, DROP COLUMN, and RENAME COLUMN
- DROP TABLE
- RENAME TABLE, TRUNCATE TABLE, and CREATE INDEX

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Self-managed SQL Server database	The SELECT permission	The SELECT permission	The permissions of the sysadmin role
ApsaraDB RDS for SQL Server instance	The read and write permissions	The read and write permissions	The read and write permissions

For more information about how to create and authorize a database account, see the following topics:

- Self-managed SQL Server database: CREATE USER
- ApsaraDB RDS for SQL Server instance: Create an account for an ApsaraDB RDS SQL Server instance

#### Process of incremental data migration

To prevent data migration failures caused by dependencies among objects, DTS migrates the schemas and data from the source SQL Server database in the following order:

- 1. Migrate the schemas of tables, views, synonyms, user-defined types, rules, defaults, and plan guides.
- 2. Perform full data migration.
- 3. Migrate the schemas of SQL stored procedures, SQL functions, triggers, and foreign keys.
- 4. Perform incremental data migration.

(?) Note During schema migration and full data migration, we recommend that you do not perform data definition language (DDL) operations on the source objects. Otherwise, the objects may fail to be migrated.

## Before you begin

Before you configure a data migration task, configure log settings on the self-managed SQL Server database.

1. Run the following command in the self-managed SQL Server database to change the recovery model to full:

```
use master;
GO
ALTER DATABASE <database_name> SET RECOVERY FULL WITH ROLLBACK IMMEDIATE;
GO
```

#### Parameters:

<database\_name>: the name of the source database. Example:

```
use master;
GO
ALTER DATABASE mytestdata SET RECOVERY FULL WITH ROLLBACK IMMEDIATE;
GO
```

2. Run the following command to create a logical backup for the source database. Skip this step if you have already created a logical backup.

BACKUP DATABASE <database\_name> TO DISK='<physical\_backup\_device\_name>'; GO

#### Parameters:

- <database\_name>: the name of the source database.
- <physical\_backup\_device\_name>: the storage path and file name of the backup file.

#### Example:

BACKUP DATABASE mytestdata TO DISK='D:\backup\dbdata.bak'; GO

3. Run the following command to back up the log entries of the source database:

```
BACKUP LOG <database_name> to DISK='<physical_backup_device_name>' WITH init;
GO
```

#### Parameters:

- <database\_name>: the name of the source database.
- <physical\_backup\_device\_name>: the storage path and file name of the backup file.

#### Example:

```
BACKUP LOG mytestdata TO DISK='D:\backup\dblog.bak' WITH init;
GO
```

4. Create a clustered index for each source table. For more information, see Create Clustered Indexes.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.

## Data Migration Migrate data from a self-managed database to Alibaba Cloud

1.Configure Source a	nd Destination	2.Configu	ure Migration Types and Objects >		3.Map name modificati	on >	4.Pre	check
*	Task Name: S	QL Server_TO_RDS						
Source Database	L							
	Instance Tunes			_				
* In	stance Region:	User-Created Databa	se with Public IP Address	•	Get ID Address Segment	of DTS		
*[	)atabase Type:	Sol Server		• •	occar Address beginning			
* Hostname	or IP Address:							
	Port Number:	1433						
* Data	abase Account:	dtstest						
* Datab	ase Password:	•••••	4	6	Test Connectivity	⊘ Passed		
Destination Database								
*	Instance Type:	RDS Instance	,	v				
* In	stance Region:	Singapore	,	¥				
* RD	S Instance ID:			•				
* Data	abase Account:	dtstest						
* Datab	ase Password:	•••••	4	6	Test Connectivity	⊘ Passed		
							Cancel Se	t Whitelist and Next
Section	Param	neter	Description					
N/A	Task I	Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.					
	Instance Type Instance Region		Select an instance type database. In this examı <b>Public IP Address</b> .	e b pl	based on the d e, select <b>User</b> -	eployme - <b>Createc</b>	nt of the sou d Database v	rce with
			<b>Note</b> If you see the network environme information, see	ele ne ee	ect other insta ent for the selt Preparation o	nce types f-manage verview.	s, you must d ed database.	eploy For
			If the instance type is s IP Address, you do no parameter.	et ot	t to <b>User-Crea</b> need to specif	ated Dat Ty the Ins	abase with tance Regio	Public n
			<b>Note</b> If a whit Server database, you the whitelist of the or Segment of DTS ne blocks of DTS servers	el la la ex	list is configure must add the ( itabase. You ca ct to <b>Instance</b>	ed for the CIDR block an click <b>G</b> Region	e self-manago ks of DTS serv <b>et IP Addres</b> to obtain the	ed SQL vers to s <b>s</b> CIDR
	Datab	ase Type	Select <b>SQLServer</b> .					

Section	Parameter	Description
Database	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed SQL Server database. In this example, enter the public IP address.
		Enter the service port number of the self-managed SQL Server database. The default port number is <b>1433</b> .
	Port Number	<b>Note</b> The service port of the self-managed SQL Server database must be accessible over the Internet.
	Dat abase Account	Enter the account that is used to log on to the self-managed SQL Server database. For information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the database account.
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
	Instance Type	Select RDS Instance.
	Instance Type Instance Region	Select <b>RDS Instance</b> . Select the region where the destination RDS instance resides.
	Instance Type Instance Region RDS Instance ID	Select RDS Instance.         Select the region where the destination RDS instance resides.         Select the ID of the destination RDS instance.
	Instance Type Instance Region RDS Instance ID Database Account	Select RDS Instance.         Select the region where the destination RDS instance resides.         Select the ID of the destination RDS instance.         Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
	Instance Type Instance Region RDS Instance ID Database Account	Select RDS Instance.         Select the region where the destination RDS instance resides.         Select the ID of the destination RDS instance.         Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
	Instance Type Instance Region RDS Instance ID Database Account	Select RDS Instance.         Select the region where the destination RDS instance resides.         Select the ID of the destination RDS instance.         Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
	Instance Type Instance Region RDS Instance ID Database Account	Select RDS Instance.         Select the region where the destination RDS instance resides.         Select the ID of the destination RDS instance.         Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
Destinatio	Instance Type Instance Region RDS Instance ID Database Account	Select RDS Instance.         Select the region where the destination RDS instance resides.         Select the ID of the destination RDS instance.         Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.

Section	Parameter	Description
		Enter the password of the database account.
	Dat abase Password	Note After you specify the information about the RDS instance, you can click Test Connectivity next to Database Password to check whether the information is valid. If the information is valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Then, modify the information based on the check results.

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### □ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

Cloud

1.Configure Source and Destination 2.Configure Migration Types and \* Migration Types: 🗹 Schema Migration 🛛 🗹 Full Data Migration 📝 Incremental Data Migration Note: Incremental data migration does not support triggers. For more information, see Reference Note: do not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database cleans up the log too early, the DTS incremental task may fail Available Selected (To edit an object name or its filter, hover over the object and click Edit.) Learn more. Q Expand the tree before you perform a glo Q 🗄 🦢 test123 📔 dtstestdata 🗉 😑 testdb > < Select All Remove All \*Rename Databases and Tables: O Not Change Database and Table Names Change Database and Table Names \* Retry Time for Failed Connection 720 Minutes 🕐 Information: 1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database. 2. Do not do DDL operation during structure and full migration, otherwise the task may fail .... Cancel Previous Save Preche

Setting	Description
Select the migratio n types	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration. In this example, select all of the three migration types.</li> </ul>
	<b>Note</b> If <b>Incremental Data Migration</b> is not selected, do not write data to the self-managed SQL Server database during full data migration. This ensures data consistency between the source and destination databases.

## Data Migration Migrate data from a self-managed database to Alibaba Cloud

Setting	Description		
	Select one or more objects from the <b>Available</b> section and click the <b>y</b> icon to move the objects to the <b>Selected</b> section.		
Select the objects to be migrate d	<ul> <li>Note</li> <li>A single data migration task can migrate incremental data from only one database. To migrate incremental data from multiple databases, you must create a data migration task for each database.</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>		
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.		
Specify the retry time for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.		
ons to the source or destinati on databas	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.		
е			

#### 8. Click Precheck.

## ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

#### 1

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the destination RDS instance.

## What to do next

The database accounts that are used for data migration have the read and write permissions. After data migration is complete, you must delete the accounts of both the self-managed SQL Server database and the ApsaraDB RDS for SQL Server instance to ensure security.

# 6.2.2. Migrate full data from a self-managed SQL Server database to an ApsaraDB RDS for SQL

## Server instance

This topic describes how to migrate full data from a self-managed SQL Server database to an ApsaraDB RDS for SQL Server instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. To migrate full data from a self-managed SQL Server database, you can select Schema Migration and Full Data Migration as the migration types.

**?** Note For information about how to migrate data without service disruptions, see Migrate incremental data from a self-managed SQL Server database to an ApsaraDB RDS for SQL Server instance.

## Prerequisites

- The version of the self-managed SQL Server database is 2005, 2008, 2008 R2, 2012, 2014, 2016, 2017, or 2019.
  - ? Note
    - If a primary/secondary switchover is performed in an SQL Server cluster or Always On availability group (AOAG), the log serial numbers of the primary database and the secondary database become inconsistent. In this case, DTS considers that the source database does not support resumable transmission, and the migration task fails. Therefore, you cannot use an SQL Server cluster or an SQL Server AOAG as the source database.
    - If you migrate data between different versions of databases, make sure that the database versions are compatible.
- The available storage space of the ApsaraDB RDS for SQL Server instance is larger than the total size of the data in the self-managed SQL Server database.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS automatically creates a destination database in the ApsaraDB RDS for SQL Server instance. However, if the name of the source database is invalid, you must create a database in the ApsaraDB RDS for SQL Server instance before you configure the data migration task.

(?) Note For more information about the naming conventions of ApsaraDB RDS for SQL Server databases and how to create a database, see Create a database on an ApsaraDB RDS for SQL Server instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

## Limits

• DTS does not migrate data of the SQL\_VARIANT type.

• DTS does not migrate the schemas of assemblies, service brokers, full-text indexes, full-text catalogs, distributed schemas, distributed functions, CLR stored procedures, CLR scalar-valued functions, CLR table-valued functions, internal tables, systems, or aggregate functions.

## Migration types

• Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, view, trigger, synonym, SQL stored procedure, SQL function, plan guide, user-defined type, rule, default, and sequence.

• Full dat a migration

DTS migrates historical data of the required objects from the self-managed SQL Server database to the destination database in the ApsaraDB RDS for SQL Server instance.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Pricing.

## Permissions required for database accounts

Database	Schema migration	Full data migration
Self-managed SQL Server database	The SELECT permission	The SELECT permission
ApsaraDB RDS for SQL Server instance	The read and write permissions	The read and write permissions

For information about how to create and authorize a database account, see the following topics:

- Self-managed SQL Server dat abase: CREATE USER
- ApsaraDB RDS for SQL Server instance: Create an account for an ApsaraDB RDS SQL Server instance

## Process of full data migration

To prevent data migration failures caused by dependencies among objects, DTS migrates the schemas and data from the source SQL Server database in the following order:

- 1. Migrate the schemas of tables, views, synonyms, user-defined types, rules, defaults, and plan guides.
- 2. Perform full data migration.
- 3. Migrate the schemas of SQL stored procedures, SQL functions, triggers, and foreign keys.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.

- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination	2.Configure Migration Types and Objects >	3.Map name modification > 4.Precheck					
* Task Name: S	Task Name: SQL Server_TO_RDS						
Source Database							
* Instance Type:	User-Created Database with Public IP Address						
* Instance Region:	Singapore	Get IP Address Segment of DTS					
* Database Type:	SQLServer v						
* Hostname or IP Address:							
* Port Number:	1433						
* Database Account:	dtstest						
* Database Password:	••••••	Test Connectivity 📀 Passed					
-							
Destination Database							
* Instance Type:	RDS Instance						
* Instance Region:	Singapore						
* RDS Instance ID:	-						
Database Account:	dtstest						
* Database Password:	••••••	Test Connectivity OPassed					
		Cancel Set Whitelist and Next					

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
		Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database with Public IP Address</b> .
	Instance Type	<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .

Section	Parameter	Description
		If the instance type is set to <b>User-Created Database with Public</b> <b>IP Address</b> , you do not need to specify the <b>Instance Region</b> parameter.
Source Database	Instance Region	Note If a whitelist is configured for the self-managed SQL Server database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
	Database Type	Select SQLServer.
	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed SQL Server database. In this example, enter the public IP address.
		Enter the service port number of the self-managed SQL Server database. The default port number is <b>1433</b> .
	Port Number	<b>Note</b> The service port of the self-managed SQL Server database must be accessible over the Internet.
	Dat abase Account	Enter the account that is used to log on to the self-managed SQL Server database. For information about the permissions that are required for the account, see Permissions required for database accounts.
	Dat abase Password	Enter the password of the database account. <b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
	Instance Type	Select RDS Instance.
	Instance Region	Select the region where the destination RDS instance resides.
	RDS Instance ID	Select the ID of the destination RDS instance.
	Dat abase Account	Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
Destinatio n		

Database Section	Parameter	Description		
		Enter the password of the database account.		
Dat abase Password	Dat abase Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

**Note** DTS adds the CIDR blocks of DTS servers to the whitelist of the destination RDS instance. This ensures that DTS servers can connect to the destination RDS instance.

#### 7. Select the migration types and the objects to be migrated.

Availab					
	le			Selected (To edit an object name or its filter, hover ov	ver the object and clic
Exna	nd the tree before you perform a	ald Q		Edit.) Learn more.	
- Cope	test123	gron i T		Q	
•	testdb			📔 dtstestdata	
			>		
			<		
Salact	AU				
Select	Ali			Remove All	
Select /	all a Databases and Tables:	<ul> <li>Do Not Change Data</li> </ul>	abase and Table Names	Remove All O Change Database and Table Names	

#### Dat a Transmission Service

Setting	Description
Select the migratio n types	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> <li>In this example, select Schema Migration and Full Data Migration.</li> <li>Note To ensure data consistency, we recommend that you do not write data to the self-managed SQL Server database during data migration.</li> </ul>
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the ) icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the



icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.

**?** Note We recommend that you do not manually stop a data migration task. Otherwise, the data migrated to the destination instance will be incomplete. You can wait until the data migration task automatically stops.

12. Switch your workloads to the destination RDS instance.

## What to do next

The database accounts that are used for data migration have the read and write permissions. After data migration is complete, you must delete the accounts of both the self-managed SQL Server database and the ApsaraDB RDS for SQL Server instance to ensure security.

## 6.3. Source database: Oracle

## 6.3.1. Migrate the schema of a self-managed

## Oracle database to a PolarDB for Oracle cluster

This topic describes how to use Advanced Database & Application Migration (ADAM) to migrate the schema of a self-managed Oracle database to a PolarDB for Oracle cluster. Then, you can use Data Transmission Service (DTS) to migrate data.

## Prerequisites

- The version of the self-managed Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.
- The self-managed Oracle database is running in ARCHIVELOG mode. Archived log files are accessible and a suitable retention period is set for archived log files. For more information, see Managing Archived Redo Log Files.
- Supplement al logging, including SUPPLEMENT AL\_LOG\_DAT A\_PK and SUPPLEMENT AL\_LOG\_DAT A\_UI, is enabled for the self-managed Oracle database. For more information, see Supplement al Logging.
- The tables to be migrated from the self-managed Oracle database contain primary keys or UNIQUE NOT NULL indexes.
- If the version of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.

• If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

- If the PolarDB for Oracle cluster is not located in the China (Beijing), China (Hangzhou), China (Zhangjiakou), China (Shenzhen), or China (Shanghai) region, you must apply a public endpoint for the PolarDB for Oracle cluster. For more information, see Apply for an endpoint.
- The polar\_comp\_redwood\_raw\_names parameter is set to *off* for the PolarDB for Oracle cluster. For more information, see Configure cluster parameters.

## Step 1: Collect the database schema

The following table describes the methods that you can use to collect the database schema by using ADAM.

Collection method	Scenario	References
Collect database information online	<ul> <li>The source Oracle database has a public IP address. You can access the database over the Internet.</li> <li>The source Oracle database has no public IP addresses. You can access the database by using Alibaba Cloud Database Gateway.</li> </ul>	Collect database information
Download Database Collector to collect database information	The source database cannot be accessed over the Internet. You must download a Database Collector client to collect the information of the database.	Collect database information

## Step 2: Analyze the database profile

- 1. Log on to the ADAM console. In the left-side navigation pane, click **Evaluate DB**.
- 2. On the Collect DB Information tab, click Next: View Source Database Profile.
- 3. Click **Create Profile**. In the Create Profile panel, set the basic information of the profile and click **Create**.

Parameter	Description
Profile Name	Specify an informative profile name for easy identification.
Туре	The value of this parameter is set to <b>Oracle</b> .

Parameter	Description
Report Language Type	Select Chinese or English.
Data File	<ul><li>i. Click Upload.</li><li>ii. Select the compressed file that is exported in Step 1: Collect the database schema.</li></ul>

4. Return to the profile list, find the database profile, and then click **Details** in the **Actions** column. The source database profile is analyzed based on the following aspects: scale, session, risk, hotspot, complexity, and load.



## Step 3: Evaluate the compatibility of the destination database

- 1. In the left-side navigation pane, click **Evaluate DB**.
- 2. Click the Evaluate Dest. DB tab.
- 3. Click Create Project . In the Create Project panel, set the basic information of the project and click Create.

Parameter	Description
Project Name	Specify an informative project name for easy identification.
Project Type	Select POLARDB-O.
Destination Database Version	Select POLARDB O.
Report Language	Select Chinese or English.

Parameter	Description	
Source Database Profile	Select the source database profile that is created in Step 2.	
Migrate LOBs to OSS	Select <b>yes</b> or <b>no</b> to specify whether to migrate large objects (LOBs) to Object Storage Service (OSS).	
Select Evaluated Schemas	<ul> <li>Select the schemas that you want to evaluate.</li> <li>i. In the Not Selected section, select the schemas that you want to evaluate.</li> <li>ii. Click the icon to add the schemas to the Selected section.</li> </ul>	

After the project is created, the state of the project changes to **Processing**.

4. Wait until the project evaluation is completed, and then click **Details** in the **Actions** column of the project. You can view the compatibility, specifications, and migration risks of the destination database. ADAM provides corresponding solutions.

## Step 4: Migrate and revise the database schema

Transformation type	Preparation	References
Online transformation	<ul> <li>For your account permissions: <ul> <li>If you are using an Alibaba Cloud account, skip this step.</li> <li>If you are using a Resource Access Management (RAM) user, perform the following steps:</li> <li>Log on to the ADAM console by using your Alibaba Cloud account.</li> <li>On the Transform &amp; Migrate DB page, grant the relevant permissions to the RAM user.</li> </ul> </li> <li>Add the CIDR blocks of ADAM servers to the whitelist of the PolarDB for Oracle cluster. For more information, see Configure the whitelist and Add CIDR blocks to the whitelist of a cluster.</li> <li>If the ADAM server that is used to migrate a database and the destination database are in the same region, you must add the private CIDR block of the ADAM server in the region to the whitelist of the destination database.</li> <li>If the ADAM server that is used to migrate a database and the destination database.</li> <li>If the ADAM server that is used to migrate a database and the destination database.</li> </ul>	Online transformation.

• Warning Online transformation implements only the migration of schemas. When you migrate the schema of a database, ADAM evaluates and verifies the objects and provides solutions for incompatible objects. You can trouble shoot issues based on the error messages to ensure better compatibility between the source and destination databases.

## What's next

Migrate data from a self-managed Oracle database to a PolarDB for Oracle cluster

## 6.3.2. Migrate data from a self-managed Oracle

## database to a PolarDB for Oracle cluster

This topic describes how to migrate data from a self-managed Oracle database to a PolarDB for Oracle cluster by using Data Transmission Service (DTS).

#### Prerequisites

The schema of the source Oracle database is migrated to the destination PolarDB for Oracle cluster by using Advanced Database & Application Migration (ADAM). For more information, see Steps 1 to 4 in Migrate the schema of a self-managed Oracle database to a PolarDB for Oracle cluster.

#### Billing

You are billed when incremental data migration is in progress, including the period when incremental data migration is paused. For more information, see Pricing.

## Precautions

- During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- If the engine version of your Oracle database is 12c or later, the names of tables to migrate cannot exceed 30 bytes in length.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

## Preparations

Log on to the self-managed Oracle database, create an account that you want to use to collect data, and then grant permissions to the account.

**?** Note If you have created a database account and the account has the permissions that are listed in the following table, skip this step.

Database	Schema migration	Full data migration	Incremental data migration
Self-managed Oracle database	The permissions of the schema owner	The permissions of the schema owner	Database administrator (DBA)
PolarDB cluster	The permissions of the schema owner	The permissions of the schema owner	The permissions of the schema owner

- Self-managed Oracle database: CREATE USER and GRANT
- PolarDB cluster: Create an account

#### Procedure

- 1. Log on to the DTS console. In the left-side navigation pane, click **Data Migration**.
- 2. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 3. In the upper-right corner of the page, click **Create Migration Task**.
- 4. Configure the source and destination databases.

* Tas	sk Name: Oracle_To_PolarDB	3-0	
Source Database			
* Inst	tance Type: User-Created D	atabase with Public IP Address	DTS support type
* Instar	nce Region: China (Hangzho	ou)	Get IP Address Segment of DTS
* Data	abase Type: Oracle		T
* Hostname or	IP Address:		
* Po	ort Number: 1521		
* Inst	tance Type:  Non-RAC Insta	ance ORAC or PDB Instance	
	* SID: testsid		
* Databa	se Account: dtstest		
* Database	abase Password:		
Destination Database			
* Inst	tance Type: PolarDB		v
* Insta	nce Region: China (Hangzho	ou)	¥
* PolarDB I	instance ID:		•
* Datal	base Name: dtstestdata		
* Databa	se Account: dtstest		
* Database	e Password: •••••		◀> Test Connectivity
			Cancel Assess Data Migration to Cloud Set Whitelist and Next
Section	Parameter	Description	
N/A Task Name that y task.		The task name tha that you specify a task. You do not n	at DTS automatically generates. We recommend a descriptive name that makes it easy to identify the need to specify a unique task name.

Section	Parameter	Description	
	Instance Type	The instance type of the source database. In this example, <b>User</b> - <b>Created Database with Public IP Address</b> is selected for this parameter.	
		<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .	
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to configure the <b>Instance Region</b> parameter.	
	Instance Region	Note If a whitelist is configured for the self-managed Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.	
	Database Type	Select <b>Oracle</b> .	
Source	Hostname or IP Address	The endpoint that is used to connect to the self-managed Oracle database.	
	Port Number	The service port number of the self-managed Oracle database. Default value: <b>1521</b> .	
		<b>Note</b> The service port of the self-managed Oracle database must be accessible over the Internet.	
Database	Instance Type	• If you select <b>Non-RAC Instance</b> , you must specify the <b>SID</b> parameter.	
		<ul> <li>If you select RAC Instance, you must specify the Service Name parameter.</li> </ul>	
	Dat abase Account	The account of the self-managed Oracle database. For information about the permissions that are required for the account, see <b>Preparations</b> .	

Section	Parameter	Description				
		The password of the database account.				
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.				
Destinatio n Database	Instance Type	Select PolarDB.				
	Instance Region	The region where the destination PolarDB for Oracle cluster resides.				
	PolarDB Instance ID	The ID of the destination PolarDB for Oracle cluster.				
	Dat abase Name	The name of the destination database.				
	Dat abase Account	The database account of the destination PolarDB cluster. For information about the permissions that are required for the account, see Preparations.				
	Dat abase Password	The password of the database account.				
		<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.				

#### 5. In the lower-right corner of the page, click Set Whitelist and Next.

**Note** DTS adds the CIDR blocks of DTS servers to the whitelist of the destination PolarDB cluster. This ensures that DTS servers can connect to the destination PolarDB cluster.

#### 6. Select the migration types and objects to migrate.

Available				Selected (To edit an obje	ct name or its filter, hove	over the object and dick
Expand th	ne tree before you perform a glo	Q				
🕀 🦢 test1	123			😑 dtstestdata		4
💽 🦳 testo	др					
			>			
			-			
			`			
Select All				Demons All		
				Charge Database and	Table Nerver	
*Rename Dat	tabases and Tables:	Do Not Change Database a	nd Table Names	U Unande Database and	able Names	
*Rename Dat * Retry Time	tabases and Tables:	Do Not Change Database a	(?)	<ul> <li>Change Database and</li> </ul>	Table Names	
*Rename Dat * Retry Time	tabases and Tables: (©	Do Not Change Database a       720       Minutes	(2)	<ul> <li>Change Database and</li> </ul>	Table Names	

<ul> <li>Schema migration is complete. Therefore, you do not need to select Schema Migration in this step. For more information, see Migrate the schema of a self-managed Oracle database to a PolarDB for Oracle cluster.</li> <li>To perform only full data migration, select only Full Data Migration.</li> <li>To ensure service continuity during data migration, select Full Data Migration and Incremental Data Migration.</li> <li>Note         <ul> <li>If Incremental Data Migration is not selected, do not write data to the self-managed Oracle database during full data migration. This ensures data consistency between the source and destination databases.</li> <li>The following SQL operations can be synchronized during incremental data migration:</li></ul></li></ul>	Setting	Description				
<ul> <li>Note</li> <li>If Incremental Data Migration is not selected, do not write data to the self-managed Oracle database during full data migration. This ensures data consistency between the source and destination databases.</li> <li>The following SQL operations can be synchronized during incremental data migration:         <ul> <li>DML operations: INSERT, UPDATE, and DELETE</li> <li>DRL</li> </ul> </li> </ul>		<ul> <li>Schema migration is complete. Therefore, you do not need to select Schema Migration in this step. For more information, see Migrate the schema of a self-managed Oracle database to a PolarDB for Oracle cluster.</li> <li>To perform only full data migration, select only Full Data Migration.</li> <li>To ensure service continuity during data migration, select Full Data Migration and Incremental Data Migration.</li> </ul>				
<ul> <li>Select</li> <li>Select the migration ntypes</li> <li>ALTER TABLE, including ADD COLUMN, ADD INDEX, DROP COLUMN, DROP INDEX, MODIFY COLUMN, and RENAME COLUMN</li> <li>DROP TABLE</li> <li>RENAME TABLE, TRUNCATE TABLE, and CREATE INDEX</li> <li>If Incremental Data Migration is selected, DTS migrates the ROWID columns from tables without primary keys. This allows you to deduplicate and verify data in tables without primary keys.</li> </ul>	Select the migratio n types	<ul> <li>Incremental Data Migration is not selected, do not write data to the selfmanaged Oracle database during full data migration. This ensures data consistency between the source and destination databases.</li> <li>The following SQL operations can be synchronized during incremental data migration:         <ul> <li>PML operations: INSERT, UPDATE, and DELETE</li> <li>DDI:</li> <li>CREATE TABLE</li> <li>Onte If a CREATE TABLE statement contains partitioning clauses, subpartitioning clauses, or functions, DTS does not migrate the CREATE TABLE AS SELECT statements.</li> <li>ALTER TABLE, including ADD COLUMN, ADD INDEX, DROP COLUMN, DROP INDEX, MOIPY COLUMN, and RENAME COLUMN</li> <li>DROP TABLE</li> <li>RENAME TABLE, TRUNCATE TABLE, and CREATE INDEX</li> </ul> </li> <li>In Incremental Data Migration is selected, DTS migrates the ROWID columns from tables without primary keys. This allows you to deduplicate and verify data in tables without primary keys.</li> </ul>				
Setting	Description					
---	--					
Select the objects that you want to migrate	Select objects from the Available section and click the Selected section. The schema of the objects has been migrated to the destination PolarDB for Oracle cluster by using ADAM. Warning In this scenario, the schema of the source database is migrated to the destination database by using ADAM. We recommend that you do not use the object name mapping feature. Otherwise, the data migration task fails.					
Specify whether to rename objects	In this scenario, the schema of the source database is migrated to the destination database by using ADAM. We recommend that you do not use the object name mapping feature. Otherwise, the data migration task fails.					
Specify the retry time range for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.					
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.					
е						

#### 7. Perform a precheck and start the data migration task.

i. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed

item to view details.

- You can troubleshoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- ii. After the data migration task passes the precheck, click Next.

- iii. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- iv. Click Buy and Start to start the data migration task.
- 8. After data migration is complete, stop the data migration task.

• Warning To minimize the negative impact of data migration on your business, we recommend that you switch your workloads to the destination PolarDB for Oracle cluster and prepare a rollback solution. You can migrate increment al data from the destination PolarDB for Oracle cluster back to the source Oracle database in real time. For more information, see Switch workloads to the destination database. If you do not need to switch your workloads, perform the following steps to stop the data migration task.

• Full data migration

Do not manually stop a task during full data migration. Otherwise, the system may fail to migrate all data. Wait until the migration task automatically ends.

- Incremental data migration
   The task does not automatically end during incremental data migration. You must manually stop the migration task.
  - a. Wait until the task progress bar shows **Incremental Data Migration** and **The migration task is not delayed**. Then, stop writing data to the source database for a few minutes. In some cases, the progress bar shows the delay time of **incremental data migration**.
  - b. After the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.

Task Name  Search by migration task name.	Search Sort: Default Sorting V Status: All V	Tag
Task ID/Name:     A Nov 2021, 09:59:01 Created	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag Completed
Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Incremental Data M 1 n The migration task is not delayed.
Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) < < 1 > >

## 6.3.3. Migrate data from a self-managed Oracle

## database to a PolarDB for Oracle cluster

PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. It is compatible with MySQL, PostgreSQL, and Oracle database engines. PolarDB provides superior performance in storage and computing to meet diverse requirements of enterprises. This topic describes how to migrate data from a self-managed Oracle database to a PolarDB for Oracle cluster by using Data Transmission Service (DTS).

## Prerequisites

- The version of the self-managed Oracle dat abase is 9i, 10g, 11g, 12c, 18c, or 19c.
- The self-managed Oracle database is running in ARCHIVELOG mode. Archived log files are accessible and a suitable retention period is set for archived log files. For more information, see Managing Archived Redo Log Files.
- Supplement al logging, including SUPPLEMENT AL\_LOG\_DAT A\_PK and SUPPLEMENT AL\_LOG\_DAT A\_UI, is enabled for the self-managed Oracle database. For more information, see Supplement al Logging.
- The tables to be migrated from the self-managed Oracle database contain primary keys or UNIQUE

NOT NULL indexes.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data from the source database will overwrite the data in the destination database after the task is resumed.
- If the version of your Oracle database is 12c or later, the names of the tables to be migrated cannot exceed 30 bytes in length.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## **Migration types**

```
Migration type Description
```

Migration type	Description			
	DTS migrates the schemas of the required objects from the source database to the destination PolarDB for Oracle cluster (PolarDB cluster in short). DTS supports schema migration for the following types of objects: table, view, synonym, trigger, stored procedure, function, package, and user-defined type.			
Schema migration	<b>Note</b> In this scenario, DTS is incompatible with triggers. If an object contains triggers, data will become inconsistent between the source and destination databases.			
	DTS migrates historical data of the required objects from the source database to the destination PolarDB cluster.			
Full data migration	<b>Note</b> During schema migration and full data migration, we recommend that you do not perform data definition language (DDL) operations on the required objects. Otherwise, the objects may fail to be migrated.			
	<ul> <li>DTS retrieves redo log files from the self-managed Oracle database. Then, DTS synchronizes incremental data from the self-managed Oracle database to the destination PolarDB cluster.</li> <li>The following SQL operations can be synchronized during incremental data migration:</li> <li>DML operations: INSERT, UPDATE, and DELETE</li> <li>DDL: <ul> <li>CREATE TABLE</li> </ul> </li> </ul>			
Incremental data migration	<b>Note</b> If a CREATE TABLE statement contains partitioning clauses, subpartitioning clauses, or functions, DTS does not migrate the statement. In addition, DTS does not migrate the CREATE TABLE AS SELECT statements.			
	<ul> <li>ALTER TABLE, including ADD COLUMN, ADD INDEX, DROP COLUMN, DROP INDEX, MODIFY COLUMN, and RENAME COLUMN</li> <li>DROP TABLE</li> <li>RENAME TABLE TRUNCATE TABLE, and CREATE INDEX</li> </ul>			
	Incremental data migration allows you to ensure service continuity when you migrate data from the self-managed Oracle database to the destination PolarDB cluster.			

## Before you begin

Log on to the source Oracle database, create an account for data collection, and grant permissions to the account.

**?** Note If you have created a database account and the account has the permissions that are listed in the following table, skip this step.

Database	Schema migration	Full data migration	Incremental data migration
Self-managed Oracle database	The permissions of the schema owner	The permissions of the schema owner	Database administrator (DBA)
PolarDB cluster	The permissions of the schema owner	The permissions of the schema owner	The permissions of the schema owner

For more information about how to create and authorize a database account, see the following topics:

- Self-managed Oracle database: CREATE USER and GRANT
- PolarDB cluster: Create an account

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

* Tasł	k Name: O	Pracle_To_PolarDB-O						
Source Database								
* Insta	ance Type:	User-Created Database	with Public IP Address	•	DTS support typ	De		
* Instan	ce Region:	China (Hangzhou)		•	Get IP Address	Segment of I	DTS	
* Data	base Type:	Oracle		•				
* Hostname or I	P Address:							
* Por	rt Number:	1521						
* Insta	ance Type:	Non-RAC Instance	RAC or PDB Instance					
	* SID:	testsid						
* Databas	e Account:	dtstest						
* Database	Password:	•••••		<b>\$</b> >	Test Conne	ectivity		
Destination Database								
* Insta	ance Type:	PolarDB		Ŧ				
* Instan	ce Region:	China (Hangzhou)		٣				
* PolarDB In	stance ID:			•				
* Datab	ase Name:	dtstestdata						
* Database Account: dtstest								
* Database Password:		•••••		<b>\$</b> >	Test Conne	ectivity		
						Cancel	Assess Data Migration to Cloud	Set Whitelist and Next
Section	Para	meter	Description					

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to use a unique task name.
	Instance Type	Select an instance type based on the deployment of the source database. In this example, select User-Created Database with Public IP Address. Note If you select other instance types, you must deploy the network environment for the source database. For more information, see Preparation overview.
	Instance Region	If you select User-Created Database with Public IP Address as the instance type, you do not need to specify the Instance Region parameter. <b>Note</b> If a whitelist is configured for the self-managed Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
	Database Type	Select Oracle.
	Database Type Hostname or IP Address	Select <b>Oracle</b> . Enter the endpoint that is used to connect to the self-managed Oracle database. In this example, enter the public IP address.
Source Dat abase	Database Type Hostname or IP Address Port Number	Select Oracle.         Enter the endpoint that is used to connect to the self-managed Oracle database. In this example, enter the public IP address.         Enter the service port number of the self-managed Oracle database. The default port number is 1521.         ⑦ Note The service port of the self-managed Oracle database must be accessible over the Internet.
Source Dat abase	Database Type Hostname or IP Address Port Number Instance Type	Select Oracle.         Enter the endpoint that is used to connect to the self-managed Oracle database. In this example, enter the public IP address.         Enter the service port number of the self-managed Oracle database. The default port number is 1521.         ⑦ Note The service port of the self-managed Oracle database must be accessible over the Internet.         • If you select Non-RAC Instance, you must specify the SID parameter.         • If you select RAC Instance, you must specify the SID parameter.
Source Dat abase	Dat abase Type Hostname or IP Address Port Number Instance Type Dat abase Account	Select Oracle.         Enter the endpoint that is used to connect to the self-managed Oracle database. In this example, enter the public IP address.         Enter the service port number of the self-managed Oracle database. The default port number is 1521.         ⑦ Note The service port of the self-managed Oracle database must be accessible over the Internet.         • If you select Non-RAC Instance, you must specify the SID parameter.         • If you select RAC Instance, you must specify the SID parameter.         • If you select RAC Instance, you must specify the SID parameter.         • If you select RAC Instance, you must specify the SID parameter.         • If you select RAC Instance, you must specify the SID parameter.         • If you select RAC Instance, you must specify the SID parameter.         • If you select RAC Instance, you must specify the Service Name parameter.

Section	Parameter	Description				
		Enter the password of the database account.				
	Database Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.				
	Instance Type	Select PolarDB.				
	Instance Region	Select the region where the destination PolarDB cluster resides.				
	PolarDB Instance ID	Select the ID of the destination PolarDB cluster.				
	Dat abase Name	Enter the name of the destination database.				
Destinatio n Database	Database Account	Enter the database account of the destination PolarDB cluster. For information about the permissions that are required for the account, see Before you begin.				
		Enter the password of the database account.				
	Database Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.				

## 6. In the lower-right corner of the page, click Set Whitelist and Next.

**Note** DTS adds the CIDR blocks of DTS servers to the whitelist of the destination PolarDB cluster. This ensures that DTS servers can connect to the destination PolarDB cluster.

## 7. Select the migration types and the objects to be migrated.

Cloud

2.Configure Migration Types and \* Migration Types: 🗹 Schema Migration 🛛 🗹 Full Data Migration 📝 Incremental Data Migration Note: Incremental data migration does not support triggers. For more information, see Reference Note: do not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database cleans up the log too early, the DTS incremental task may fail Available Selected (To edit an object name or its filter, hover over the object and click Edit.) Learn more. Q Expand the tree before you perform a glo Q 🗄 🦢 test123 📔 dtstestdata 🗉 😑 testdb > < Select All Remove All \*Rename Databases and Tables: O Not Change Database and Table Names O Change Database and Table Names \* Retry Time for Failed Connection 720 Minutes 🕐 Information: 1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database. 2. Do not do DDL operation during structure and full migration, otherwise the task may fail Cancel Previous Save Prechec

Setting	Description
	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> </ul>
Select the migratio n types	<ul> <li>Note</li> <li>If Incremental Data Migration is not selected, we recommend that you do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.</li> <li>If Incremental Data Migration is selected, DTS migrates the ROWID columns in tables without primary keys. This allows you to deduplicate and verify data in tables without primary keys.</li> </ul>

Setting	Description					
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>					
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.  Note If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.					
Specify the retry time for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.					
e						

#### 8. In the lower-right corner of the page, click **Precheck**.

## ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the jicon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.

- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the destination PolarDB cluster.

## 6.3.4. Migrate data from a self-managed Oracle

## database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a self-managed Oracle database to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a self-managed Oracle database, you can select all of the supported migration types to ensure service continuity.

## Prerequisites

- The version of the self-managed Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.
- Supplement al logging, SUPPLEMENTAL\_LOG\_DATA\_PK, and SUPPLEMENTAL\_LOG\_DATA\_UI are enabled for the self-managed Oracle database. For more information, see Supplement al Logging.
- The self-managed Oracle database runs in ARCHIVELOG mode. Archived log files of the Oracle database are accessible, and an appropriate retention period is specified for the archived log files. For more information, see Managing Archived Redo Log Files.
- The size of available storage in the RDS instance is larger than the size of data that you want to migrate from the self-managed Oracle database.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

• Table names in the RDS instance are not case-sensitive. If a table name in the self-managed Oracle database contains uppercase letters, ApsaraDB RDS for MySQL converts all uppercase letters to lowercase letters and then creates the table.

If the self-managed Oracle database contains identical table names that differ only in capitalization, the table names are identified as duplicates. As a result, the "The object already exists" message may be displayed during schema migration. To prevent table name conflicts in the RDS instance, you can rename the migrated objects by using the object name mapping feature of DTS. For more information, see Object name mapping.

• DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. If the name of the source database is invalid, you must manually create a destination database in the RDS instance before you configure a data migration task. For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## **Migration** types

• Schema migration

DTS supports schema migration for tables and indexes. DTS does not support schema migration for the following types of objects: views, synonyms, triggers, stored procedures, stored functions, packages, and user-defined data types. DTS has the following limits on schema migration for tables and indexes:

- DTS does not support schema migration for nested tables. DTS converts clustered tables and index-organized tables into standard tables in the destination database.
- DTS does not support schema migration for function-based indexes, domain indexes, bit map indexes, or reverse indexes.
- Full data migration

DTS migrates the historical data of specified objects from the self-managed Oracle database to the destination database in the RDS instance.

• Incremental data migration

DTS uses the round-robin algorithm to retrieve redo log files from the self-managed Oracle database. Then, DTS synchronizes incremental data from the self-managed Oracle database to the destination database in the RDS instance. Incremental data migration ensures service continuity when you migrate data from the self-managed Oracle database to the destination database in the RDS instance.

# SQL operations that can be synchronized during incremental data migration

- INSERT, DELETE, and UPDATE
- CREATE TABLE

**?** Note DTS cannot synchronize the CREATE TABLE operations that are performed to create tables in which functions are nested.

- ALT ER TABLE, ADD COLUMN, DROP COLUMN, RENAME COLUMN, and ADD INDEX
- DROP TABLE
- RENAME TABLE, TRUNCATE TABLE, and CREATE INDEX

## Data type mappings

For more information, see Data type mappings between heterogeneous databases.

## Before you begin

Log on to the self-managed Oracle database, create an account that you want to use to collect data, and grant permissions to the account.

(?) **Note** If you created an account that is granted the permissions listed in the following table, you can skip this step.

Database	Schema migration	Full data migration	Incremental data migration
Self-managed Oracle database	Permissions of the schema owner	Permissions of the schema owner	DBA

Database	Schema migration	Full data migration	Incremental data migration
ApsaraDB RDS for MySQL instance	Write permissions on the destination database	Write permissions on the destination database	Write permissions on the destination database

For more information about how to create an account and grant permissions to the account, see the following topics:

- Self-managed Oracle database: CREATE USER and GRANT.
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account for an ApsaraDB RDS for MySQL instance.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the self-managed Oracle database and RDS instance.

1.Configure Source and Des	stination D	atabases 🔰 2.Config	gure Migration Types and Objects					
* Task	Name: C	racle_TO_RDS		]				
Source Database								
* Insta	nce Type:	User-Created Database wi	th Public TD Addrass	•	DTS support type			
* Instanc	o Pogion	Chine (Use set au)	arrable ir Addess	-	Got ID Address Segment of	DTC		
Instanc	e Region:	China (Hangzhou)		•	Get IP Address Segment of	015		
* Datab	ase Type:	Oracle		۳				
* Hostname or IP	Address:							
* Port	t Number:	1521						
* Insta	nce Type:	Non-RAC Instance RA	AC Instance					
	* SID:							
* Database	Account:			٦				
* Database F	Password:	•••••		\$>	Test Connectivity			
				_				
Destination Database								
* Insta	nce Type:	RDS Instance		٣				
* Instanc	e Region:	China (Hangzhou)		۳				
* RDS Ins	stance ID:			•				
* Database	e Account:							
* Database F	Password:	•••••		<b>\$</b> >	Test Connectivity			
* Encryption:   Non-encrypted  SSL-encrypted								
						Cancel	Assess Data Migration t	to Cloud Set Whitelist and Next
Section	Par	ameter	Description					

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify a name that can help you identify the task. You do not need to specify a unique task name.
		Select an instance type based on the deployment of the self- managed Oracle database. In this example, select <b>User-Created</b> <b>Database with Public IP Address</b> .
	Instance Type	<b>Note</b> If you select other instance types, you must set up the environment that is required for the self-managed Oracle database. For more information, see <b>Preparation overview</b> .
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to configure the <b>Instance Region</b> parameter.
	Instance Region	<b>Note</b> If an IP address whitelist is configured for the self- managed Oracle database, you must add the CIDR blocks of DTS servers to the IP address whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
	Database Type	Select <b>Oracle</b> .
	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed Oracle database. In this example, enter the public IP address.
	Port Number	Enter the port number that is used to connect to the self-managed Oracle database. The port must be accessible over the Internet. The default port number is <b>1521</b> .
Source Dat <i>a</i> base	Instance Type	<ul> <li>If you select Non-RAC Instance, you must configure the SID parameter.</li> <li>If you select RAC or PDB Instance, you must configure the Service Name parameter.</li> </ul>
	Dat abase Account	Enter the account that you created in the self-managed Oracle database. For more information about the permissions that are required for the account, see Before you begin.

Section	Parameter	Description			
	Dat abase Password	Enter the password of the preceding account. <b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			
	Instance Type	Select RDS Instance.			
	Instance Region	Select the region where the RDS instance resides.			
	RDS Instance ID	Select the ID of the RDS instance.			
Destinatio	Database Account	Enter the account that is used to connect to the RDS instance. For more information about the permissions that are required for the account, see Before you begin.			
Database		Enter the password of the preceding account.			
	Dat abase Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			

6. In the lower-right corner of the page, click Set Whitelist and Next.

## 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects that you want to migrate.

1.Configure So	urce and Destination > 2.Configure Migration Types and > 3.Advanced Settings > 4.Precheck
* Migratio triggers. Fo Note: do cleans u	on Types: Schema Migration Full Data Migration Incremental Data Migration Note: Incremental data migration does not support or more information, see Reference o not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database o the log too early, the DTS incremental task may fail
Available Expand Expand Expand Expand Expand Expand Select All *Rename I * Retry Tir Informati 1. Data mining the source of the sou	Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Edit) Laam more. Selected (To edit an object name or its filter, hover over the object and dick Selected (To edit an object name or its filter, hover over the object and ble hames Selected (To edit an object name or its filter, hover over the object and selecter of the object name or its filter over the object name or its filter over over the object name or its filter over over object name or its filter over over over over over over over ov
Setting	Description
Select the migratio n types	<ul> <li>If you want to perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>If you want to ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> <li>Note If you do not select Incremental Data Migration, make sure that no data is written to the self-managed Oracle database during full data migration. This ensures data consistency between the self-managed Oracle database and the RDS instance.</li> </ul>

Setting	Description
Select the objects that you want to migrate	Select one or more objects in the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
	<ul> <li>Note</li> <li>You can select columns, tables, or databases.</li> <li>By default, the name of an object that is migrated to the RDS instance remains the same as that in the self-managed Oracle database. You can use the object name mapping feature to rename the objects that are migrated to the RDS instance. For more information, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the RDS instance. For more information, see Object name mapping.
Specify the retry time range for failed	By default, if DTS fails to connect to the self-managed Oracle database or RDS instance, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the self-managed Oracle database and RDS instance within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
connecti ons to the self- manage d Oracle databas e or RDS	<b>Note</b> Within the time range in which DTS attempts to reconnect to the self- managed Oracle database and RDS instance, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at the earliest opportunity after the self-managed Oracle database and RDS instance are released.
instance	

## 8. Click Precheck.

## ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.

<sup>(1)</sup> 

- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

5	sk Name	Search Sort: Default Sorting V Status: All V	S Tag
	Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate	Task Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
	4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed Incremental Data M 1 n The migration task is not delayed.
~	3 Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) $\langle \langle 1 \rangle \rangle$

12. Switch your workloads over to the RDS instance.

## What to do next

The accounts that are used to migrate data are granted the read and write permissions. After the data migration is complete, you must delete the accounts of the self-managed Oracle database and the RDS instance to ensure database security.

## Additional information

DTS supports reverse data transmission when you migrate data from a self-managed Oracle database to an ApsaraDB RDS for MySQL instance. You can use this feature to synchronize data changes from the ApsaraDB RDS for MySQL instance to the self-managed Oracle database. If you want to use the reverse data transmission feature, submit a ticket.

## 6.3.5. Migrate data from a self-managed Oracle

## database to a PolarDB for MySQL cluster

This topic describes how to migrate data from a self-managed Oracle database to a PolarDB for MySQL cluster by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a self-managed Oracle database, you can select all of the supported migration types to ensure service continuity.

## Prerequisites

• The version of the self-managed Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.

- Supplemental logging, including SUPPLEMENTAL\_LOG\_DATA\_PK and SUPPLEMENTAL\_LOG\_DATA\_UI, is enabled for the self-managed Oracle database. For more information, see Supplemental Logging.
- The self-managed Oracle database is running in ARCHIVELOG mode. Archived log files are accessible and a suitable retention period is set for archived log files. For more information, see Managing Archived Redo Log Files.
- The available storage space of the PolarDB for MySQL cluster is larger than the total size of the data in the self-managed Oracle database.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination cluster, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination cluster after the task is resumed.
- Table names in the PolarDB for MySQL cluster are case-insensitive. If a table name in the source Oracle database contains uppercase letters, PolarDB for MySQL converts all uppercase letters to lowercase letters before creating the table.
   If the source Oracle database contains identical table names that differ only in capitalization, these table names are identified as duplicate. During scheme migration, the following message is returned:

table names are identified as duplicate. During schema migration, the following message is returned: "The object already exists." To prevent name conflicts in the destination database, you can rename the migrated objects by using the object name mapping feature. For more information, see Object name mapping.

• DTS automatically creates a destination database in the PolarDB for MySQL cluster. However, if the name of the source database is invalid, you must manually create a database in the PolarDB for MySQL cluster before you configure the data migration task. For more information about how to create a database and the database naming conventions, see Create a database.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## Migration types

• Schema migration

DTS supports schema migration for tables and indexes. DTS does not support schema migration for the following types of objects: view, synonym, trigger, stored procedure, function, package, and user-defined type. DTS has the following limits on schema migration for tables and indexes:

- DTS does not support schema migration for nested tables. DTS converts clustered tables and index-organized tables (IOTs) into common tables in the destination database.
- DTS does not support schema migration for function-based indexes, domain indexes, bit map indexes, or reverse indexes.
- Full data migration

DTS migrates historical data of the required objects from the self-managed Oracle database to the destination database in the PolarDB for MySQL cluster.

• Incremental data migration

DTS retrieves redo log files from the self-managed Oracle database. Then, DTS synchronizes incremental data from the self-managed Oracle database to the destination database in the PolarDB for MySQL cluster. Incremental data migration allows you to ensure service continuity when you migrate data from the self-managed Oracle database to the destination database.

## Data type conversion

Oracle and PolarDB for MySQL are heterogeneous databases. DTS does not ensure that the schemas of the source and destination databases are consistent after schema migration. We recommend that you evaluate the impact of data type conversion on your business. For more information, see Data type mappings between heterogeneous databases.

## Before you begin

Log on to the source Oracle database, create an account for data collection, and grant permissions to the account.

**Note** If you have created a database account and the account has the permissions that are listed in the following table, skip this step.

Database	Schema migration	Full data migration	Incremental data migration
Self-managed Oracle database	The permissions of the schema owner	The permissions of the schema owner	The database administrator (DBA) permission
PolarDB for MySQL cluster	The permissions of the schema owner	The permissions of the schema owner	The permissions of the schema owner

For more information about how to create and authorize a database account, see the following topics:

- Self-managed Oracle database: CREATE USER and GRANT
- PolarDB for MySQL cluster: Create a database account and Manage database accounts for a cluster

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.

Source Database					
*	Instance Type:	User-Created Data	base in ECS Instance	~	DTS support type
* In	stance Region:	China (Hangzhou)		$\sim$	
<b>*</b> EC	CS Instance ID:	1 Automation (1)	Per l	•	
*[	Database Type:	Oracle		$\sim$	
		ADAM evaluates the	compatibility and transformation plan of Oracle mi	gration	to other databases for you free of charge Go to View Demo
	* Port Number:	1521			
*	Instance Type:	Non-RAC Instance	e ORAC or PDB Instance		
	* SID:	-			
* Data	abase Account:	100 m			
* Datab	base Password:	•••••		<b>\$</b> >	Test Connectivity
-					
Destination Database					
*	Instance Type:	PolarDB		~	
* To	ctanco Rogion-	China (Usasahaw)			
	* Instance Region: China (Hangzhou			~	
↑ PolarL	* PolarDB Instance ID:			•	
* Data	* Database Account:				
* Datab	base Password:	•••••	9	Þ	Test Connectivity
Section	Parame	eter	Description		
N/A	Task Name		DTS automatically general specify an informative na use a unique task name.	ates me	s a task name. We recommend that you to identify the task. You do not need to
			Select an instance type b database. In this example <b>Public IP Address</b> .	ase e, se	ed on the deployment of the source elect <b>User-Created Database with</b>
	Instance Type		<b>Note</b> If you sele the network environme information, see <b>Prepa</b>	ect o ent i rati	other instance types, you must deploy for the source database. For more <mark>on overview</mark> .

Section	Parameter	Description		
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to specify the <b>Instance Region</b> parameter.		
	lnst ance Region	Note If a whitelist is configured for the self-managed Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.		
	Database Type	Select Oracle.		
Source Database	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed Oracle database. In this example, enter the public IP address.		
	Port Number	Enter the service port number of the self-managed Oracle database. The port must be accessible over the Internet. The default port number is <b>1521</b> .		
	<ul> <li>If you select Non-RAC Instance, you must specify the SID parameter.</li> <li>If you select RAC Instance, you must specify the Service Nan parameter.</li> </ul>			
	Dat abase Account	Enter the account of the self-managed Oracle database. For information about the permissions that are required for the account, see Before you begin.		
		Enter the password of the database account.		
	Dat abase Password	<b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.		
	Instance Type	Select PolarDB.		
	Instance Region	Select the region where the destination PolarDB for MySQL cluster resides.		
	PolarDB Instance ID	Select the ID of the destination PolarDB for MySQL cluster.		
Destinatio	Dat abase Account	Enter the database account of the destination PolarDB for MySQL cluster. For more information about the permissions that are required for the account, see Before you begin.		

Betabase	Parameter	Description
	Enter the password of the database account.	
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.

## 6. In the lower-right corner of the page, click Set Whitelist and Next.

(?) Note In this step, DTS adds the CIDR blocks of DTS servers to the whitelist of the destination PolarDB for MySQL cluster. This ensures that DTS servers can connect to the destination cluster.

#### 7. Select the migration types and the objects to be migrated.

Note: do not clean up the incremental data log generated by the source cleans up the log too early, the DTS incremental task may fail Available Expand the tree before you perform a glol Q	database after the	DTS task is started when the DTS full task is running. If the source database Selected (To edit an object name or its filter, hover over the object and clicl
Available Expand the tree before you perform a gloi Q		Selected (To edit an object name or its filter, hover over the object and click
Expand the tree before you perform a glod Q		Edit.) Learn more.
🖃 📂 test123		Q
		ei dtstestdata
e estab		
	>	
	<	
Jelec All		Remove All
*Rename Databases and Tables: <ul> <li>Do Not Change Database a</li> </ul>	and Table Names	Change Database and Table Names
* Retry Time for Failed Connection 720 Minutes	0	
Information: 1. Data migration only copies the data and schema in the source database i in the source database. 2. Do not do DDL operation during structure and full migration, otherwise t	and saves the copy the task may fail	in the destination database. The process does not affect any data or schema

Setting	Description			
Select the migratio n types	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> <li>Note If Incremental Data Migration is not selected, do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.</li> </ul>			
	Select one or more objects from the <b>Available</b> section and click the <b>&gt;</b> icon to move the objects to the <b>Selected</b> section.			
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination cluster, the name of the object remains the same as that in the self-managed Oracle database. You can use the object name mapping feature to rename the objects that are migrated to the destination cluster. For more information, see Object name mapping.</li> </ul>			
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination cluster. For more information, see Object name mapping.			
Specify the retry time for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.			
the source or destinati on	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.			
6				

8. In the lower-right corner of the page, click **Precheck**.

## ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.

	Task Name   Search by migration task name.	Search Sort: Default Sorting V Status: All	▼ ● Tag
	Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Dup	olicate Task Upgrade Configure Monitoring and Alerting Change password Edit Tag
	4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Incremental Data & 1 n The migration task is not delayed.
E	Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) < < 1 > >

12. Switch your workloads to the destination PolarDB for MySQL cluster.

## What to do next

The database accounts that are used for data migration have the read and write permissions. After data migration is complete, you must delete the accounts of both the self-managed Oracle database and the PolarDB for MySQL cluster to ensure database security.

## More information

> Document Version: 20220712

DTS supports reverse data transmission when you migrate data from a self-managed Oracle database to a PolarDB for MySQL cluster. You can use this feature to synchronize data changes from the PolarDB for MySQL cluster to the self-managed Oracle database. To do this, submit a ticket.

# 6.3.6. Migrate data from a self-managed Oracle database to a PolarDB-X 1.0 instance

This topic describes how to migrate data from a self-managed Oracle database to a PolarDB-X 1.0 instance by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration. When you migrate data from a self-managed Oracle database, you can select the two migration types to ensure service continuity.

## Prerequisites

- The version number of the self-managed Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.
- Supplement al logging, SUPPLEMENTAL\_LOG\_DATA\_PK, and SUPPLEMENTAL\_LOG\_DATA\_UI are enabled for the self-managed Oracle database. For more information, see Supplement al Logging.
- The self-managed Oracle database is running in ARCHIVELOG mode. Archived log files are accessible and a suitable retention period is set for archived log files. For more information, see Managing Archived Redo Log Files.
- The service port of the self-managed Oracle database is accessible over the Internet.
- The databases in the PolarDB-X 1.0 instance must be created based on ApsaraDB RDS for MySQL instances. DTS does not support PolarDB-X 1.0 databases that are created based on PolarDB for MySQL clusters.
- The available storage space of the PolarDB-X 1.0 instance is larger than the total size of the data in the self-managed Oracle database.

## Precautions

• DTS does not support schema migration from a self-managed Oracle database to a PolarDB-X 1.0 instance.

**?** Note During schema migration, DTS migrates the schemas of required objects, such as tables, from the source database to the destination database.

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node

to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## **Migration** types

• Full data migration

DTS migrates the historical data of required objects from the source Oracle database to the destination PolarDB-X instance.

Onte To ensure data consistency, we recommend that you do not write data to the selfmanaged Oracle database during full data migration.

• Incremental data migration

DTS retrieves redo log files from the self-managed Oracle database. Then, DTS synchronizes incremental data from the self-managed Oracle database to the destination PolarDB-X 1.0 instance. Incremental data migration allows you to ensure service continuity when you migrate data from an Oracle database to a PolarDB-X 1.0 instance.

**Note** The following SQL operations can be synchronized during incremental data migration: INSERT, DELETE, and UPDATE operations. DDL operations cannot be synchronized during incremental data migration.

## Preparations

1. Create databases and tables in the destination PolarDB-X 1.0 instance based on the tables to migrate from the self-managed Oracle database. Fore more information, see Step 2: Create a database and Step 3: Create a table.

**?** Note The data types of Oracle databases and PolarDB-X 1.0 instances do not have oneto-one correspondence. You must define the corresponding data types in PolarDB-X 1.0 instances. For more information, see Data type mappings between heterogeneous databases.

2. Log on to the source Oracle database. Then, create an account for data collection and grant

#### permissions to the account.

Database	Full data migration	Incremental data migration
Self-managed Oracle database	Permissions of the schema owner	DBA
PolarDB-X	Write permissions on the destination database	Write permissions on the destination database

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the **Migration Tasks** page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure source and destination databases.

		atabases 2.Conne	jure Migration Types and Objects	2	3.Map name modification	>	4	I.Precheck
* Task Name: Oracle_TO_DRDS								
Source Database	Source Database							
* Instan	ice Type:	User-Created Database wit	h Public IP Address	Ŧ				
* Instance	e Region:	Singapore <b>*</b>		Ŧ	Get IP Address Segment of DTS			
* Databa	ase Type:	Oracle	,	Ŧ				
* Hostname or IP	Address:							
* Port	Number:	1521						
* Instan	ice Type:	Non-RAC Instance RA	C Instance					
	* SID:	1 - 1 - 1 - 1 - 1 - 1 - 1 - 1 - 1 - 1 -						
* Database	Account:	dtstest						
* Database P	assword:	•••••	٩	Þ	Test Connectivity			
Destination Database								
Instance Type: DRDS Instance		•						
* Instance	e Region:	Singapore		۳				
* DRDS Inst	tance ID:	-		Ŧ				
* Databas	se Name:	mytest 🔻						
* Database	Account:	dtstest						
* Database P	assword:	•••••	٩	Þ	Test Connectivity			
							Cancel	Set Whitelist and Next
Section	Par	ameter	Description					
N/A Task Name That you specify a task. You do not n		ha a : n	t DTS automatically genera descriptive name that mak eed to specify a unique tas	ates. W es it ea sk nam	/e reco asy to i e.	mmend dentify the		

Section	Parameter	Description
		The instance type of the source database. In this example, User- Created Database with Public IP Address is selected.
	Instance Type	<b>Note</b> If you select other instance types, you must set up the environment that is required for the self-managed Oracle database. For more information, see <b>Preparation overview</b> .
		If you select <b>User-Created Database with Public IP Address</b> for the instance type, you do not need to configure the <b>Instance Region</b> parameter.
	Instance Region	Note If a whitelist is configured for the self-managed Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
	Database Type	Select <b>Oracle</b> .
	Hostname or IP Address	The endpoint that is used to connect to the self-managed Oracle database. In this example, the public IP address of the database is used.
	Port Number	The service port number of the self-managed Oracle database. Default value: <b>1521</b> .
	Instance Type	<ul> <li>If you select Non-RAC Instance, you must specify the SID parameter.</li> <li>If you select RAC Instance, you must specify the Service Name parameter.</li> </ul>
Source Database	Dat abase Account	The account of the self-managed Oracle database. For information about the permissions that are required for the account, see <b>Preparations</b> .

Parameter	Description		
	The password of the database account.		
Dat abase Password	Note After you specify the information about the self- managed Oracle database, you can click Test Connectivity next to Database Password to check whether the information is valid. If the information is valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Then, modify the information based on the check results.		
Instance Type	Select DRDS Instance.		
Instance Region	The region where the destination PolarDB-X 1.0 instance resides.		
DRDS Instance ID	The destination PolarDB-X 1.0 instance ID.		
Database Account	The database account of the destination PolarDB-X 1.0 instance. For information about the permissions that are required for the account, see Preparations.		
	The password of the database account.		
Dat abase Password	Note After you specify the information about the RDS instance, you can click Test Connectivity next to Database Password to check whether the information is valid. If the information is valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Then, modify the information based on the check results.		
	ParameterDat abase PasswordInstance TypeInstance RegionDRDS Instance IDDat abase AccountDat abase Password		

6. In the lower-right corner of the page, click Set Whitelist and Next.

## 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects that you want to migrate.

1.Configure So	urce and Destination > 2.Configure Migration Types and > 3.Advanced Settings > 4.Precheck	
* Migratio triggers, Fo Note: do cleans u	on Types: Schema Migration Full Data Migration Incremental Data Migration Note: Incremental data migration does not support or more information, see Reference o not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database o p the log too early, the DTS incremental task may fail	
Available Expand Expand The te The te Select All *Rename I * Retry Tir Informati 1. Data mij in the sour 2. Do not co	Selected (10 edit an object name or its filter, hover over the object and dok     the tree before you perform a glol ( )     sti223   etdb     Image: tree before you perform a glol ( )	
Setting	Description	
<ul> <li>Select the migration full data migration, select Schema Migration and Data Migration.</li> <li>If you want to ensure service continuity during data migration, select Schema Migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> <li>If you do not select Incremental Data Migration, make sure the data is written to the self-managed Oracle database during full data migratic ensures data consistency between the self-managed Oracle database and the instance.</li> </ul>		

Setting	Description			
	Select one or more objects in the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.			
Select the objects that you want to migrate	<ul> <li>Note</li> <li>You can select columns, tables, or databases.</li> <li>By default, the name of an object that is migrated to the RDS instance remains the same as that in the self-managed Oracle database. You can use the object name mapping feature to rename the objects that are migrated to the RDS instance. For more information, see Object name mapping.</li> </ul>			
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the RDS instance. For more information, see Object name mapping.			
Specify the retry time range for failed	By default, if DTS fails to connect to the self-managed Oracle database or RDS instance, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the self-managed Oracle database and RDS instance within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.			
connecti ons to the self- manage d Oracle databas e or RDS	<b>Note</b> Within the time range in which DTS attempts to reconnect to the self- managed Oracle database and RDS instance, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at the earliest opportunity after the self-managed Oracle database and RDS instance are released.			
Instance				

## 8. Click Precheck.

## ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.

<sup>(1)</sup> 

- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

6	Task Name   Search by migration task name.	Search Sort: Default Sorting	▼ Tag
	Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Dup	plicate Task Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
	4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed
	Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s)

12. Switch the workloads from the source database to the PolarDB-X 1.0 instance.

## What to do next

The database accounts that are used for data migration have the read and write permissions. After data migration is complete, you must delete the accounts of both the self-managed Oracle database and the PolarDB-X 1.0 instance to ensure security.

## 6.3.7. Migrate data from a self-managed Oracle

## database to an AnalyticDB for PostgreSQL

## instance

This topic describes how to migrate data from a self-managed Oracle database to an AnalyticDB for PostgreSQL instance by using Data Transmission Service (DTS).

## Prerequisites

- The version of the self-managed Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.
- The self-managed Oracle database is running in ARCHIVELOG mode. Archived log files are accessible and a suitable retention period is set for archived log files. For more information, see Managing Archived Redo Log Files.
- Supplement al logging, including SUPPLEMENTAL\_LOG\_DATA\_PK and SUPPLEMENTAL\_LOG\_DATA\_UI, is enabled for the self-managed Oracle dat abase. For more information, see Supplement al Logging.

• An instance is created. For more information, see Create an AnalyticDB for PostgreSQL instance.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## Precautions

- During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination database after the task is resumed.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

## **Migration types**

Migration type

Description
Migration type	Description			
Schema migration	DTS migrates the schemas of the required objects from the source database to the destination database. DTS supports schema migration for the following types of objects: table, index, constraint, function, sequence, and view.			
	<ul> <li>Oracle and AnalyticDB for PostgreSQL are heterogeneous databases. DTS does not ensure that the schemas of the source and destination databases are consistent after schema migration. We recommend that you evaluate the impact of data type conversion on your business. For more information, see Data type mappings between heterogeneous databases.</li> <li>For partitioned tables, DTS discards the partition definitions. You must define partitions in the destination database.</li> </ul>			
	DTS migrates historical data of the required objects from the source database to the destination database.			
Full data migration	<b>Note</b> During schema migration and full data migration, we recommend that you do not perform data definition language (DDL) operations on the required objects. Otherwise, the objects may fail to be migrated.			
	After full data migration is complete. DTC retrieves rede los files from the source			
Incremental data migration	<ul> <li>After full data migration is complete, DTS retrieves redo log files from the source Oracle database. Then, DTS migrates incremental data from the source Oracle database to the destination database in real time.</li> <li>DTS can synchronize the following SQL operations during incremental data migration:</li> <li>DML operations: INSERT, UPDATE, and DELETE</li> <li>DDL operation: ADD COLUMN</li> <li>Incremental data migration allows you to ensure service continuity when you migrate data from a self-managed Oracle database.</li> </ul>			

# Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Self-managed Oracle database	The permissions of the schema owner	The permissions of the schema owner	The database administrator (DBA) permission
	The read and write permissions on the destination database	The read and write permissions on the destination database	The read and write permissions on the destination database

For more information about how to create and authorize a database account, see the following topics:

- Self-managed Oracle database: CREATE USER and GRANT
- : Configure an account

Notice If you need to migrate incremental data from an Oracle database but the DBA permission cannot be granted to the database account, you can grant fine-grained permissions to the account. The following sample statements show you how to grant specific permissions to an Oracle database account.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

* Tas	k Name:	Oracle_To_ADB for PG					
	No	ote: If you need to carry out	t incremental data migration for a long t	ime,	it is recommended to use the data sy	nchronization function, which has bet	tter network stability and
richer functions. Click to buy	richer functions. Click to buy data synchronization Click to see the difference between data migration and synchronization						
Source Database							
* Inst	ance Type:	User-Created Database	in ECS Instance	~	DTS support type		
# Technol	co Degion						
Instan	ice Region	China (Hangzhou)		~			
* ECS Ir	istance ID:			•			
* Data	base Type:	Oracle		~			
		ADAM evaluates the comp	atibility and transformation plan of Orac	de m	igration to other databases for you fre	ee of charge Go to View Demo	
* Po	rt Number:	1521					
* Inst	ance Type:	■ Non-RAC Instance O	RAC or PDB Instance	_			
	* SID:	testsid					
* Databas	e Account:	dtstest					
* Database	Password	•••••	<	⊅	Test Connectivity		
Destination Database							
* Instance Type: AnalyticDB for PostgreSQL		QL	~				
* Instan	ice Region:	China (Hangzhou)		~			
* Ir	nstance ID:	gp-		•			
* Datab	ase Name:	dtstestdata					
* Databas	se Account:	dtstest					
* Database	Password		(	<b>\$</b> >	Test Connectivity		
					Cancel	Assess Data Migration to Cloud	Set Whitelist and Next
Section	Dar	amotor	Description				
Section	Pdl	ameter	Description				
			DTS automatically	n	enerates a task nan	ne Werecommend	that you
N/A	Tas	k Name	specify an informa	er et iv	e name to identify	the task. You do n	ot need to
			use a unique task	na	me.		

Section	Parameter	Description
		Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database in ECS Instance</b> .
	Instance Type	<b>Note</b> If you select other instance types, you must prepare the environment that is required for the source database. For more information, see <b>Preparation overview</b> .
	Instance Region	Select the region of the Elastic Compute Service (ECS) instance on which the self-managed Oracle database is hosted.
	ECS Instance ID	Select the ID of the ECS instance on which the self-managed Oracle database is hosted.
	Database Type	Select Oracle.
	Port Number	Enter the service port number of the self-managed Oracle database. The default port number is <b>1521</b> .
Source Dat abase	Instance Type	<ul> <li>If you select Non-RAC Instance, you must specify the SID parameter.</li> <li>RAC Instance: If you select this option, you must specify the Service Name.</li> <li>In this example, select Non-RAC Instance.</li> </ul>
	SID	Enter the system ID (SID) of the self-managed Oracle database.
	Dat abase Account	Enter the account of the self-managed Oracle database. For more information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the database account.
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
	Instance Type	Select AnalyticDB for PostgreSQL.
	Instance Region	Select the region where the destination instance resides.
	Instance ID	Select the ID of the instance.

Destinatio Gection Database	Parameter	Description
	Dat abase Name	Enter the name of the destination database.
	Dat abase Account	Enter the database account of the destination instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.
	Dat abase Password	Enter the password of the database account.

- 6. In the lower-right corner of the page, click Set Whitelist and Next. DTS adds the CIDR blocks of DTS servers to the inbound rule of the ECS instance and the whitelist of the instance. This ensures that DTS servers can connect to the source and destination instances.
- 7. Select the migration types, the operation types, and the objects to be migrated.

the DTS full task is running. If the source database bject name or its filter, hover over the object and di Q abase (SObjects)
the DTS full task is running. If the source database bject name or its filter, hover over the object and di Q abase (50bjects)
bject name or its filter, hover over the object and die Q abase (SObjects)
bject name or its filter, hover over the object and di Q abase (50bjects)
abase (50bjects)
Q abase (50bjects)
abase (SObjects)
and Table Names
nu table Names
se. The process does not affect any data or schema

Setting	Description
Select the migratio n types	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> <li>Note If Incremental Data Migration is not selected, do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.</li> </ul>
Select the operatio n types	Select the types of operations that you want to synchronize during incremental data migration. All operation types are selected by default.
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the  icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or schemas as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
e	

Setting	Description	
Enclose object names in quotatio n marks	Specify whether you need to enclose object names in quotation marks. If you select <b>Yes</b> and the following conditions are met, DTS encloses object names in single or double quotation marks during schema migration and incremental data migration.	
	• The business environment of the source database is case-sensitive but the database name contains both uppercase and lowercase letters.	
	• A source table name does not start with a letter and contains characters other than letters, digits, and special characters.	
	<b>Note</b> A source table name can contain only the following special characters: underscores (_), number signs (#), and dollar signs (\$).	
	• The names of the schemas, tables, or columns that you want to migrate are keywords, reserved keywords, or invalid characters in the destination database.	

8. Specify the primary key columns and distribution keys of the tables that you want to migrate to the instance.

			3.Advance	ed Settings	
Schema name	Table Name	Type(All) 👻	Primary Key Column	Distribution key	Definition Status(All) 👻
dtstest	ORACLETESTTABLE	Hash distrib	ID	ID	Defined
Enter a table nam	ne. Search			Total: 1 item(s), Per Page	$: 20 \text{ item(s)}  \ll  <  1  >  \gg$
				Cancel	Previous Save Precheck

#### ? Note

- For more information about primary key columns and distribution keys, see Define constraints and Define table distribution.
- If DTS identifies tables without primary keys, the option Set Primary Keys and Distribution Keys of All Tables Without Primary Keys to ROWID is displayed on the preceding page. If you select this option, DTS adds the ROWID field as the primary key and distribution key to the destination tables.

#### 9. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.

- 10. After the task passes the precheck, click Next.
- 11. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 12. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



# 6.3.8. Migrate data from a self-managed Oracle database to a Message Queue for Apache Kafka

# instance

You can use Data Transmission Service (DTS) to migrate data from a self-managed Oracle database to a Message Queue for Apache Kafka instance or a self-managed Kafka cluster. The data migration feature allows you to extend message processing capabilities. This topic describes how to migrate data from a self-managed Oracle database to a Message Queue for Apache Kafka instance.

# Prerequisites

- The version number of the self-managed Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.
- Supplement al logging, SUPPLEMENTAL\_LOG\_DATA\_PK, and SUPPLEMENTAL\_LOG\_DATA\_UI are enabled for the self-managed Oracle database. For more information, see Supplement al Logging.
- The self-managed Oracle database is running in ARCHIVELOG mode. Archived log files are accessible and a suitable retention period is set for archived log files. For more information, see Managing Archived Redo Log Files.
- The network environment is deployed for the source Oracle database. For more information, see

#### Preparation overview.

- The tables to migrate from the self-managed Oracle database contain primary keys or UNIQUE NOT NULL indexes.
- The version number of the Message Queue for Apache Kafka instance is 0.10.1.0 to 2.x. The version number of the self-managed Kafka cluster is 0.10.1.0 to 2.7.0.
- The available storage space of the destination Message Queue for Apache Kafka instance is larger than the total size of the data in the self-managed Oracle database.
- In the destination Message Queue for Apache Kafka instance, a topic is created to receive the synchronized data. For more information, see Create a topic.

# Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data from the source database overwrites the data in the destination database after the task is resumed.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

- If the version of your Oracle database is 12c or later, the names of the tables to migrate cannot exceed 30 bytes in length.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# Migration types

Migration type	Description		
Schema migration	DTS migrates the schemas of required objects from the source database to the destination database. In this scenario, DTS can migrate only the schemas of tables.		
	DTS migrates the historical data of required objects from the source database to the destination database.		
Full data migration	<b>Note</b> During schema migration and full data migration, do not perform DDL operations on the objects to migrate. Otherwise, the objects may fail to be migrated.		
Incremental data migration	After full data migration is complete, DTS retrieves redo log files from the source Oracle database. Then, DTS migrates incremental data from the source Oracle database to the destination database in real time. Incremental data migration ensures service continuity when you migrate data between self-managed databases. During incremental data migration, DTS can synchronize DML and DDL operations.		

# Preparations

Log on to the self-managed Oracle database, create an account that you want to use to collect data, and grant permissions to the account.

**Note** If you have created a database account and the account has permissions that are listed in the following table, skip this step.

Database	Schema migration	Full data migration	Incremental data migration
Self-managed Oracle database	Permissions of the schema owner	Permissions of the schema owner	Permissions of the database administrator

For more information about how to create an account and grant permissions to the account, see the following topics:

Self-managed Oracle databases: CREATE USER and GRANT

# Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of **Migration Tasks** page, select the region where the destination instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

#### Dat a Transmission Service

	* Instance Type:	User-Cr	eated Database with Public IP Address	Supported Databases			
	* Instance Region:	China (ł	Hangzhou) 🗸	Get IP Address Segment of DTS			
* Database Type: Oracle			~				
charge Go to Vie	w Demo	ADAM eva	luates the compatibility and transformation plan of Oracle migra	tion to other databases for you free of			
* Hos	stname or IP Address:			]			
	* Port Number:	1521					
	* Instance Type:		AC Instance O RAC or PDB Instance				
	* SID:			]			
	* Database Account:	dtstest		]			
•	Database Password:		4	Test Connectivity			
_							
Destination Database	2						
	* Instance Type:	User-Cr	eated Database Connected Over Express Connect, VPN Ga $$				
	* Instance Region:	China (ł	Hangzhou)	Guide			
	* Peer VPC:	vpc-					
	* Database Type:	Kafka	~				
	* IP Address:						
	* Port Number:	9092					
	Database Account:			Optional			
	Database Password:		4	Optional			
	* Topic:		~	Get Toplic list			
		Click Get	Topic List and then select the specific topic.				
	Topic for storing DDL:	Click Get	Topic List and then select the specific topic.	Get Toplic list			
	* Kafka Version	1.0	~				
	* Encryption:	Non-er	crypted OSCRAM-SHA-256				
* Whether to use K	afka schema registry:	® № C	) Yes				
Section	Parameter		Description				
N/A	Task Name		The task name that DTS automatical that you specify a descriptive name t task. You do not need to specify a u	ly generates. We recomm hat makes it easy to ident nique task name.	end tify the		
Instance Type			The access method of the source database. In this example, <b>User-</b> <b>Created Database with Public IP Address</b> is selected.				
		pe	<b>Note</b> If the source database access methods, you must set up required for the self-managed Ora information, see <b>Preparation overv</b>	e is connected over other the environment that is cle database. For more <mark>iew</mark> .			

Section	Parameter	Description
		If you select <b>User-Created Database with Public IP Address</b> for the Instance Type parameter, you do not need to configure the <b>Instance Region</b> parameter.
	Instance Region	Note If a whitelist is configured for the self-managed Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
	Database Type	Select Oracle.
	Hostname or IP Address	The IP address that is used to connect to the self-managed Oracle database. In this example, the public IP address is entered.
		The service port number of the self-managed Oracle database. Default value: <b>1521</b> .
Source Database	Port Number	<b>Note</b> The service port of the self-managed Oracle database must be accessible over the Internet.
	Instance Type	<ul> <li>If you select Non-RAC Instance, you must configure the SID parameter.</li> <li>If you select RAC or PDB Instance, you must configure the Service Name parameter.</li> </ul>
	Database Account	The account of the self-managed Oracle database. For information about the permissions that are required for the account, see <b>Preparations</b> .
		The password of the database account.
	Database Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.

Section	Parameter	Description
		Select User-Created Database Connected over Express Connect, VPN Gateway, or Smart Access Gateway.
	Instance Type	<b>Note</b> You cannot select Message Queue for Apache Kafka for the Instance Type parameter. You can use Message Queue for Apache Kafka as a self-managed Kafka database to configure data synchronization.
	Instance Region	The region where the destination Message Queue for Apache Kafka instance resides.
		The ID of the virtual private cloud (VPC) to which the destination Message Queue for Apache Kafka instance belongs. To obtain the VPC ID, you can log on to the Message Queue for Apache Kafka console and go to the Instance Details page of the Message Queue for Apache Kafka instance. On the <b>Instance Details</b> page, you can view the VPC ID.
	Peer VPC	Configuration Information         VPC ID       vpc       vsw         Zone       Zone H       Secondary Zone       None ●         Topics       50       Groups       100         Partitions       400       Billing Method       Pay-as-you-go         Created At       Dec 27, 2019, 14:42:04       Message Retention       72 Hours < Edit
	Database Type	Select Kafka.
	IP address	Enter an IP address that is included in the <b>Default Endpoint</b> parameter of the Message Queue for Apache Kafka instance. <b>Note</b> To obtain an IP address, you can log on to the Message Queue for Apache Kafka console and go to the Instance Details page of the Message Queue for Apache Kafka instance. On the <b>Instance Details</b> page, you can obtain an IP address from the <b>Default Endpoint</b> parameter.
Destinatio n Database	Port Number	The service port number of the Message Queue for Apache Kafka instance. Default value: 9092.

Section	Parameter	Description				
		The username that is used to log on to the Message Queue for Apache Kafka instance.				
	Dat abase Account	<b>Note</b> If the Message Queue for Apache Kafka instance is of the VPC Instance type, you do not need to specify the database account or database password.				
	Dat abase Password	The password of the username.				
	Торіс	Click <b>Get Topic List</b> and select a topic name from the drop-down list.				
Topic That Stores DDL Information		Click <b>Get Topic List</b> , and select a topic name from the drop-down list. The topic is used to store the DDL information. If you do not specify this parameter, the DDL information is stored in the topic that is specified by the <b>Topic</b> parameter.				
	Kafka Version	The version of the Message Queue for Apache Kafka instance.				
	Encryption	Select <b>Non-encrypted</b> or <b>SCRAM-SHA-256</b> based on your business and security requirements.				
	Whether to Use Kafka Schema Registry	<ul> <li>Kafka Schema Registry provides a serving layer for your metadata. It provides a REST ful API to store and retrieve your Avro schemas.</li> <li>No: Kafka Schema Registry is not used.</li> <li>Yes: Kafka Schema Registry is used. In this case, you must enter the URL or IP address that is registered in Kafka Schema Registry for your Avro schemas.</li> </ul>				

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types, the migration policy, and the objects to migrate.

1.Configure Source a	nd Destination	2.Configure Migration Typ	es and	3.Advanced Settings	<u> </u>	4.Precheck
<ul> <li>Migration Types</li> <li>triggers. For more i</li> </ul>	<ul> <li>Schema Migration</li> <li>see Reference</li> </ul>	<ul> <li>Full Data Migration</li> </ul>	<ul> <li>Incremental</li> </ul>	Data Migration Note: Incre	mental data migrati	on does not support
Note: do not clea cleans up the log	an up the incremental da too early, the DTS incr	ata log generated by the source emental task may fail	e database after the	e DTS task is started when the D	)TS full task is runni	ng. If the source database
Data migration a between Apsara For long-term da	pplies to short-term mig Stack databases. ta synchronization in re	ration scenarios. Typical scena al time, use the data synchroni	rios include migrati zation feature.	ng data to the doud, scaling and	d sharding database	s, and migrating data
Available				Selected (To edit an object	name or its filter, h	over over the object and dick
Expand the tree	before you perform a g	lo Q		Edit.) Learn more.		
+ 🦢	A					Q
+				dtstest0415		
+						
			>			
+ 💼			<			
+ 🦢						
	T					
+ <u>-</u>						
+						
+						
Select All						
				Remove All		
*Rename Database	s and Tables:	Do Not Change Database	and Table Names	O Change Database and T	able Names	
* Retry Time for Fa	iled Connection	720 Minutes	0			
*Source table DMS want to copy the te the target database	ONLINE_ Do you mporary table to during DDL:	🔾 Yes 💿 No				
Information:						
<ol> <li>Data migration o in the source datability</li> <li>Do not do DDL o</li> </ol>	nly copies the data and ase. peration during structur	schema in the source database e and full migration, otherwise	and saves the cop the task may fail	y in the destination database. Ti	ne process does not	affect any data or schema
				Cance	I Previous	Save Precheck
etting D	escription					
S	elect <b>Schem</b> a	a Migration, Full	Data Mig	ration, and Incre	mental Da	ta Migration.
elect	2					
ne nigratio	Vou do not v	If Incremental write data to the s	Data Mig	<b>ration</b> is not select abase during full d	cted, we re lata migrati	commend that on. This ensures

n types	data consistency between the source and destination databases.
Select the data format used in Kafka	The data that is migrated to the Kafka cluster is stored in the Avro format. You must parse the migrated data based on the Avro schema. For more information, see DTS Avro schema.

Setting	Description
Select the policy for migratin g data to Kafka partition s	Select a migration policy based on your business requirements. For more information, see Specify the policy for migrating data to Kafka partitions.
Select	Select one or more tables from the <b>Available</b> section and click the <b>&gt;</b> icon to add the
the objects	tables to the <b>Selected</b> section.
that you want to migrate	<ul> <li>Note DTS maps the table names to the topic name that you select in Step 5.</li> <li>For information about how to change the topic name, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source or destination database within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
ons to the source or destinati on	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
databas e	

8. In the lower-right corner of the page, click **Precheck**.

#### ♥ Notice

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- After you troubleshoot the issues based on the causes, you can run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

### Stop the migration task

**Warning** We recommend that you prepare a rollback solution to migrate incremental data from the destination database to the source database in real time. This allows you to minimize the negative impact of switching your workloads to the destination database. For more information, see Switch workloads to the destination database. If you do not need to switch your workloads, you can perform the following steps to stop the migration task.

• Full data migration

Do not manually stop a task during full data migration. Otherwise, the system may fail to migrate all data. Wait until the migration task automatically ends.

• Incremental data migration

The task does not automatically end during incremental data migration. You must manually stop the migration task.

- i. Wait until the task progress bar shows **Incremental Data Migration** and **The migration task is not delayed**. Then, stop writing data to the source database for a few minutes. In some cases, the progress bar shows the delay time of **incremental data migration**.
- ii. After the status of incremental data migration changes to The migration task is not delayed, manually stop the migration task.

	sk Name 🗸 Search by migration task name.	Search Sort: Default Sorting V Status: All V	fag
	Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
	4 Nov 2021, 09:59:01 Created Schema Migration <b>100%</b>	Full Data Migration 100%(Migrated Rows: 1)	Completed Incremental Data M 1 n The migration task is not delayed.
~	3 Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) <

#### What's next

The database accounts that are used for data migration have read and write permissions. After data migration is complete, you must delete the account of the self-managed Oracle database. You must also modify the permissions of the Resource Access Management (RAM) user in the destination Kafka

instance. For more information, see Grant permissions to RAM users.

# 6.4. Source database: PostgreSQL

# 6.4.1. Migrate incremental data from a self-managed PostgreSQL database (version 10.1 to13) to an ApsaraDB RDS for PostgreSQL instance

You can use Data Transmission Service (DTS) to migrate incremental data between PostgreSQL databases. The source or destination database can be a self-managed PostgreSQL database or an ApsaraDB RDS for PostgreSQL instance. DTS supports schema migration, full data migration, and incremental data migration. You can create a task that includes all three migration types to ensure service continuity. This topic describes how to migrate incremental data from a self-managed PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance.

# Prerequisites

- The self-managed PostgreSQL database is of a version from 10.1 to 13.0.
- An ApsaraDB RDS for Post greSQL instance is created. For more information, see Create an ApsaraDB RDS for Post greSQL instance.

(?) Note To ensure compatibility, you must make sure that the database version of the ApsaraDB RDS for PostgreSQL instance is the same as the version of the self-managed PostgreSQL database.

• The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total size of the data in the self-managed PostgreSQL database.

# Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- During incremental data migration, if you select a schema to migrate, take note of the following requirements: If you create a table in the schema or run the RENAME command to rename a table in the schema, you must run the ALTER TABLE schema.table REPLICA IDENTITY FULL; command before you write data to the table.

(?) Note Replace the scheme and table in the preceding sample command with the actual scheme name and table name.

- To ensure that the accurate delay time of incremental data migration is displayed, DTS adds a heartbeat table named *dts\_postgres\_heartbeat* to the source database.
- During incremental data migration, DTS creates a replication slot in the source database. The replication slot is prefixed with <a href="https://dts\_sync">dts\_sync</a>. DTS automatically clears historical replication slots every 90 minutes to reduce storage usage.

**?** Note If the data migration task is released or fails, DTS automatically clears the replication slot. If a primary/secondary switchover is performed on the ApsaraDB RDS for PostgreSQL instance, you must log on to the secondary database to clear the replication slot manually.

Qu	ery Editor	Query History	Query History					Scratch Pad				×				
1	<pre>SELECT * FROM pg_replication_slots;</pre>															
Da	ta Output	Explain Mes	sag	ges Noti	fica	ations										
	slot_name name			plugin name		slot_type text	datoid oid	2	database name	temporary boolean		active boolean		active_pid integer		-
1	dts_sync_of	าน		pgoutput		logical	16		dtstestdata	false		true			-	I

• To ensure that the data migration task runs as expected, you can perform a primary/secondary switchover only on a V11 ApsaraDB RDS for PostgreSQL instance. In this case, you must set the rds\_failover\_slot\_mode parameter to sync . For more information, see Logical Replication Slot Failover.

Q Warning If you perform a primary/secondary switchover on a self-managed PostgreSQL database or an ApsaraDB RDS for PostgreSQL instance of a version other than V11, the data migration task stops.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

# Limits

- A single data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each source database.
- The name of the source database cannot contain hyphens (-), for example, dts-testdata.
- If a primary/secondary switchover is performed on the source database during incremental data migration, the transmission cannot be resumed.
- Data may be inconsistent between the primary and secondary nodes of the source database due to synchronization delay. Therefore, you must use the primary node as the data source when you migrate data.

**?** Note We recommend that you migrate data during off-peak hours. You can modify the transfer rate of full data migration based on the read/write performance of the source database. For more information, see Modify the transfer rate of full data migration.

- Incremental data migration does not support the BIT data type.
- During increment al data migration, DTS migrates only data manipulation language (DML) operations. DML operations include INSERT, DELETE, and UPDATE.

② Note Only data migration tasks that are created after October 1, 2020 can migrate data definition language (DDL) operations. You must create a trigger and function in the source database to obtain the DDL information before you configure the task. For more information, see Use triggers and functions to implement incremental DDL migration for PostgreSQL databases.

- After your workloads are switched to the destination database, newly written sequences do not increment from the maximum value of the sequences in the source database. Therefore, you must query the maximum value of the sequences in the source database before you switch your workloads to the destination database. Then, you must specify the queried maximum value as the starting value of the sequences in the destination database.
- DTS does not check the validity of metadata such as sequences. You must manually check the validity of metadata.

# Billing

Migration type	Task configuration fee	Internet traffic fee		
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from		
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.		

# Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Self-managed PostgreSQL database	The USAGE permission on pg_catalog	The SELECT permission on the objects to be migrated	The permissions of the superuser role
ApsaraDB RDS for PostgreSQL instance	The CREATE and USAGE permissions on the objects to be migrated	The permissions of the schema owner	The permissions of the schema owner

For more information about how to create a database account and grant permissions to the account, see the following topics:

- Self-managed PostgreSQL database: CREATE USER and GRANT
- ApsaraDB RDS for PostgreSQL instance: Create an account on an ApsaraDB RDS for PPAS instance.

# Data migration process

The following table describes how DTS migrates the schemas and data of the source PostgreSQL database. The process prevents data migration failures that are caused by dependencies between objects.

**Note** For more information about schema migration, full data migration, and incremental data migration, see **Terms**.

Data migration process	Description
	DTS migrates the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates to the destination database.
1. Schema migration	<b>Note</b> DTS does not migrate plug-ins. In addition, DTS does not migrate functions that are written in the C programming language.
2. Full data migration	DTS migrates all historical data of the required objects to the destination database.
3. Schema migration	DTS migrates the schemas of triggers and foreign keys to the destination database.
	DTS migrates incremental data of the required objects to the destination database. Incremental data migration ensures service continuity of self-managed applications.
4. Incremental data migration	<ul> <li>Note</li> <li>During incremental data migration, DTS migrates only data manipulation language (DML) operations. DML operations include INSERT, DELETE, and UPDATE.</li> <li>Incremental data migration does not support the BIT data type.</li> </ul>

# Preparation

- 1. Log on to the server where the self-managed PostgreSQL database resides.
- 2. Set the value of the wal\_level parameter in the *postgresql.conf* configuration file to logical

·		
# - Settin	ngs -	
wal_level	= logical	<pre># minimal, replica, or logical # (change requires restart)</pre>
⑦ Note	Skip this step if you c	lo not need to perform incremental data migration.

3. Add the CIDR blocks of DTS servers to the *pg\_hba.conf* configuration file of the self-managed PostgreSQL database. Add only the CIDR blocks of the DTS servers that reside in the same region as the destination database. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

**Note** For more information, see The pg\_hba.conf File. Skip this step if you have set the IP address in the pg\_hba.conf file to 0.0.0/0.

-				
# "loca	al" is	for Unix domain	socket connections only	
local	all	all		peer
# IPv4	local	connections:		
host	all	all	0.0.0/0	trust
# IPv6	local	connections:		
host	all	all	::1/128	md5

- 4. (Optional) Create a trigger and function in the source database to obtain the DDL information. For more information, see Use triggers and functions to implement incremental DDL migration for PostgreSQL databases.
  - Onte Skip this step if you do not need to migrate DDL operations.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and De	sunation L	Jatabases 2.Conng						
* Task	: Name: 👔	PostgreSQL_TO_RDS						
Source Database								
* Instan	nce Type:	User-Created Database wit	h Public IP Address	٣				
* Instance	e Region:	Singapore		•	Get IP Address Se	gment of DTS		
* Databa	ase Type:	PostgreSQL		•				
* Hostname or IP	Address:							
* Port	Number:	5432		]				
* Databas	se Name:	mytestdata						
* Database	Account:	dtstest						
* Database P	assword:	*****		Ф	Test Connec	tivity 🔗 Passed		
Destination Database								
* Insta	ince Type:	RDS Instance		•				
* Instanc	ce Region:	Singapore		•				
* RDS Ins	stance ID:			•				
* Databa	ase Name:	mytestdata						
* Database	e Account:	dtstest						
* Database	Password:	•••••	٩	6	Test Connectivity	⊘ Passed		
								Cancel Set Whitelist and Next
Section	Par	rameter	Description					
N/A	Ta	sk Name	DTS automatical specify a name t does not need to	ly ge hat o b be	enerates a t can help yo unique.	ask name. V u identify th	Ve recon e task. T	nmend that you he task name
	-							

Section	Parameter	Description
	Instance Type	Select an instance type based on the deployment of the self- managed database. In this example, select <b>User-Created Database</b> with Public IP Address.
		<b>Note</b> If you select other instance types, you must set up the environment that is required for the self-managed database. For more information, see <b>Preparation overview</b> .
	Instance Region	If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to specify the <b>Instance Region</b> parameter.
	Database Type	Select PostgreSQL.
	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed PostgreSQL database. In this example, enter the public IP address.
Source Database	Port Number	Enter the service port number of the self-managed PostgreSQL database. The port must be accessible over the Internet.
	Dat abase Name	Enter the name of the self-managed PostgreSQL database.
	Dat abase Account	Enter the account that is used to log on to the self-managed PostgreSQL database. For information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the database account.
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
	Instance Type	Select RDS Instance.
	Instance Region	Select the region where the RDS instance resides.
	RDS Instance ID	Select the ID of the destination RDS instance.

Section	Parameter	Description		
	Dat abase Name	Enter the name of the destination database in the RDS instance. The name can be different from the name of the source database.		
Destinatio n Database		<b>Note</b> Before you configure the data migration task, make sure that the destination database exists in the ApsaraDB RDS for PostgreSQL instance. If the destination database does not exist, you need to create the database first. For more information, see Create a database.		
Database Enter the Account see Perm		Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.		
		Enter the password of the database account.		
	Dat abase Password	<b>?</b> Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### **Warning**

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.

7. Select the migration type and the objects you want to migrate.

1.Configure Sou	rce and Destination 2.Configure Migration Types and 3.Advanced Settings A.Precheck
* Migratic triggers. Fo Note: do cleans u	Types: Schema Migration 🗹 Full Data Migration 🔽 Incremental Data Migration Note: Incremental data migration does not support more information, see Reference not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database the log too early, the DTS incremental task may fail
Available Expand Expand Select All *Rename E * Retry Tin Informati 1. Data mig in the sourd 2. Do not d	Selected (To adit an object name or its filter, hover over the object and dick     it123   idb      idb
	Cancel Previous Save Precheck
Setting	Description
Select the migratio n types	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration. In this example, select all the three migration types.</li> <li>Note If Incremental Data Migration is not selected, do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.</li> </ul>

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
Select the objects that you want to migrate	<ul> <li>Note</li> <li>You can select columns, tables, or schemas as the objects to be migrated.</li> <li>After an object is migrated to the RDS instance, the name of the object remains the same as that in the self-managed PostgreSQL database. You can use the object name mapping feature to rename the objects that have been migrated to the destination RDS instance. For more information, see Object name mapping.</li> <li>If you renamed an object by using the object name mapping feature, objects that are dependent on the renamed object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the RDS instance. For more information, see Object name mapping.
Specify the window period to retry for failed connecti ons to the self- manage d databas e or RDS instance	By default, if DTS fails to connect to the self-managed database or RDS instance, DTS retries within the following 12 hours. You can also specify a window period to retry the connection based on your business requirements. If DTS is reconnected to the self-managed database and RDS instance within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS attempts to reconnect to the self-managed database or RDS instance, you are charged for using the DTS instance. We recommend that you specify a window period to retry the connection based on your business requirements. You can also release the DTS instance at the earliest opportunity after the self-managed database and RDS instance are released.

8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

#### Stop the migration task

**Warning** We recommend that you prepare a rollback solution to migrate incremental data from the destination database to the source database in real time. This allows you to minimize the negative impact of switching your workloads to the destination database. For more information, see Switch workloads to the destination database. If you do not need to switch your workloads, you can perform the following steps to stop the migration task.

• Full data migration

Do not manually stop a task during full data migration. Otherwise, the system may fail to migrate all data. Wait until the migration task automatically ends.

Incremental data migration

The task does not automatically end during incremental data migration. You must manually stop the migration task.

- i. Wait until the task progress bar shows **Incremental Data Migration** and **The migration task is not delayed**. Then, stop writing data to the source database for a few minutes. In some cases, the progress bar shows the delay time of **incremental data migration**.
- ii. After the status of incremental data migration changes to The migration task is not delayed, manually stop the migration task.

5	Task Name   Search by migration task name.	Search Sort: Default Sorting V Status: All V	Tag
	Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
	4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed Incremental Data M 1 n The migration task is not delayed.
	Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) < < 1 > >

# 6.4.2. Migrate incremental data from a selfmanaged PostgreSQL database (in PostgreSQL 10.0 or an earlier version) to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to migrate incremental data from a self-managed PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a self-managed PostgreSQL database to Alibaba Cloud, you can select all of the supported migration types to ensure service continuity.

This topic uses a self-managed database with a public IP address as an example to describe how to configure an incremental data migration task. For information about how to perform full data migration, see Migrate full data from a self-managed PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance.

# Prerequisites

- The self-managed database runs PostgreSQL 9.4.8 or a later version, PostgreSQL 9.5, PostgreSQL 9.6, or PostgreSQL10.0.
- The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total size of the data in the self-managed PostgreSQL database.
- The service port of the self-managed PostgreSQL database is accessible over the Internet.

# Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- During incremental data migration, if you select a schema to migrate, take note of the following requirements: If you create a table in the schema or run the RENAME command to rename a table in the schema, you must run the ALTER TABLE schema.table REPLICA IDENTITY FULL; command before you write data to the table.

**Note** Replace the schema and table in the preceding sample command with the actual schema name and table name.

- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

### Limits

- A single data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each source database.
- The name of the source database cannot contain hyphens (-), for example, dts-testdata.
- If a primary/secondary switchover is performed on the source database during incremental data migration, the transmission cannot be resumed.
- Data may be inconsistent between the primary and secondary nodes of the source database due to synchronization delay. Therefore, you must use the primary node as the data source when you migrate data.

(?) Note We recommend that you migrate data during off-peak hours. You can modify the transfer rate of full data migration based on the read/write performance of the source database. For more information, see Modify the transfer rate of full data migration.

- DTS does not check the validity of metadata such as sequences. You must manually check the validity of metadata.
- After your workloads are switched to the destination database, newly written sequences do not increment from the maximum value of the sequences in the source database. Therefore, you must query the maximum value of the sequences in the source database before you switch your workloads to the destination database. Then, you must specify the queried maximum value as the starting value of the sequences in the database. You can execute the following statements to query the maximum value of the sequences in the source database:

```
do language plpgsql $$
declare
  nsp name;
  rel name;
  val int8;
begin
   for nsp,rel in select nspname,relname from pg_class t2 , pg_namespace t3 where t2.relna
mespace=t3.oid and t2.relkind='S'
  loop
     execute format($_$select last_value from %I.%I$_$, nsp, rel) into val;
     raise notice '%',
     format($_$select setval('%I.%I'::regclass, %s);$_$, nsp, rel, val+1);
   end loop;
end;
$$;
```

• To ensure that the data migration task runs as expected, you can perform a primary/secondary switchover only on a V11 ApsaraDB RDS for PostgreSQL instance. In this case, you must set the rds\_failover\_slot\_mode parameter to sync . For more information, see Logical Replication Slot Failover.

Q Warning If you perform a primary/secondary switchover on a self-managed PostgreSQL database or an ApsaraDB RDS for PostgreSQL instance of a version other than V11, the data migration task stops.

• During increment al dat a migration, DTS migrates only dat a manipulation language (DML) operations. DML operations include INSERT, DELETE, and UPDATE.

**Note** Only data migration tasks that are created after October 1, 2020 can migrate data definition language (DDL) operations. You must create a trigger and function in the source database to obtain the DDL information before you configure the task. For more information, see Use triggers and functions to implement incremental DDL migration for PostgreSQL databases.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Self-managed PostgreSQL database	The USAGE permission on pg_catalog	The SELECT permission on the objects to be migrated	The permissions of the superuser role
ApsaraDB RDS for PostgreSQL instance	The CREATE and USAGE permissions on the objects to be migrated	The permissions of the schema owner	The permissions of the schema owner

For more information about how to create a database account and grant permissions to the account, see the following topics:

- Self-managed PostgreSQL database: CREATE USER and GRANT
- ApsaraDB RDS for PostgreSQL instance: Create an account on an ApsaraDB RDS for PPAS instance.

# Data migration process

The following table describes how DTS migrates the schemas and data of the source PostgreSQL database. The process prevents data migration failures that are caused by dependencies between objects.

**Note** For more information about schema migration, full data migration, and incremental data migration, see Terms.

Description
DTS migrates the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates to the destination database.
<b>Note</b> DTS does not migrate plug-ins. In addition, DTS does not migrate functions that are written in the C programming language.

Data migration process	Description
2. Full data migration	DTS migrates all historical data of the required objects to the destination database.
3. Schema migration	DTS migrates the schemas of triggers and foreign keys to the destination database.
	DTS migrates incremental data of the required objects to the destination database. Incremental data migration ensures service continuity of self-managed applications.
4. Incremental data migration	<ul> <li>Note</li> <li>During incremental data migration, DTS migrates only data manipulation language (DML) operations. DML operations include INSERT, DELETE, and UPDATE.</li> <li>Incremental data migration does not support the BIT data type.</li> </ul>

# Before you begin

- 1. Download the PostgreSQL source code from the official website, and compile and install the source code.
  - i. Download the source code from the PostgreSQL official website based on the version of the self-managed PostgreSQL database.
  - ii. Run the ./configure , make , and make install commands to in sequence to configure, compile, and install the source code.

#### 🗘 Notice

- When you compile and install PostgreSQL, the operating system version of PostgreSQL must be consistent with the GNU Compiler Collection (GCC) version.
- If an error occurs when you run the ./configure command, you can adjust the command based on the error message. For example, if the error message is readline library not found. Use --without-readline to disable readline support. , you can change the command to ./configure --without-readline .
- If you use other methods to install PostgreSQL, you must compile the ali\_decoding plug-in in a test environment that has the same OS version and GCC version.
- 2. Download the ali\_decoding plug-in provided by DTS, and compile and install the plug-in.
  - i. Download ali\_decoding.

ii. Copy the ali\_decoding directory to the contrib directory of PostgreSQL (compiled and installed).

-rw-rr	1	1107	1107	384	9月	27	2016 aclocal.m4
d rwx rwx rwx	2	1107	1107	4096	9月	27	2016 config
-rw-rr	1	root	root	374806	9月	7	10:10 config.log
-rwxr-xr-x	1	root	root	39032	9月	7	10:10 config.status
-rwxr-xr-x	1	1107	1107	471157	9月	27	2016 configure
-rw-rr	1	1107	1107	75195	9月	27	2016 configure, in
drwxrwxrwx	56	1107	1107	4096	9月	7	10:23 contrib
-rw-rr	1	1107	1107	1192	9月	27	2016 COPYRIGHT
drwxrwxrwx	3	1107	1107	4096	9月	27	2016 doc
-rw-rr	1	root	root	3638	9月	7	10:10 GNUmakefile
-rw-rr	1	1107	1107	3638	9月	27	2016 GNUmakefile.in
-rw-rr	1	1107	1107	283	9月	27	2016 HISTORY
-rw-rr	1	1107	1107	75065	9月	27	2016 INSTALL
-rw-rr	1	1107	1107	1489	9月	27	2016 Makefile
-rw-rr	1	1107	1107	1209	9月	27	2016 README
drwxrwxrwx	16	1107	1107	4096	9月	7	10:10 src

iii. Go to the ali\_decoding directory and replace the content of the Makefile file with the following script:

```
# contrib/ali decoding/Makefile
MODULE big = ali decoding
MODULES = ali_decoding
OBJS
      = ali decoding.o
DATA = ali decoding--0.0.1.sql ali decoding--unpackaged--0.0.1.sql
EXTENSION = ali decoding
NAME = ali decoding
#subdir = contrib/ali decoding
#top_builddir = ../..
#include $(top builddir)/src/Makefile.global
#include $(top srcdir)/contrib/contrib-global.mk
#PG CONFIG = /usr/pgsql-9.6/bin/pg config
#pgsql_lib_dir := $(shell $(PG_CONFIG) --libdir)
#PGXS := $(shell $(PG CONFIG) --pgxs)
#include $(PGXS)
# Run the following commands to install the source code.
ifdef USE PGXS
PG CONFIG = pg config
PGXS := $(shell $(PG CONFIG) --pgxs)
include $(PGXS)
else
subdir = contrib/ali decoding
top builddir = ../..
include $(top_builddir)/src/Makefile.global
include $(top srcdir)/contrib/contrib-global.mk
endif
```

iv. Go to the ali\_decoding directory, run the make and make install commands in sequence to compile ali\_decoding and obtain the files required to install ali\_decoding.

v. Copy the following files to the specified location.



3. Create a database and schema in the destination instance based on the database and schema information of the objects to be migrated. The schema name of the source and destination databases must be the same. For more information, see Create a database on an ApsaraDB RDS for PostgreSQL instance and Appendix: User and schema management.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the **Migration Tasks** page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.

1.Configure Source and De	stination Databases	2.Configure Migration Types and Objects	$\geq$	3.Map name modification	$\rightarrow$	4.Precheck	
* Task	Name: PostgreSQL_TO_R	DS	]				
Source Database							
* Instan	nce Type: User-Created D	atabase with Public IP Address	•				
* Instance	e Region: Singapore		•	Get IP Address Segment of DTS			
* Databa	ase Type: PostgreSQL		٣				
* Hostname or IP	Address:						
* Port	Number: 5432						
* Databas	se Name: mytestdata						
* Database	Account: dtstest						
* Database P	assword:		<⊅	Test Connectivity 🛛 Passed			
Destination Database							
* Instance Type: RDS Instance			•				
* Instance Region: Singapore			•				
* RDS Instance ID:		i ginakan	•				
* Database Name: mytestdata			_				
* Database Account: dtstest		-	<u></u>	Test Connectivity @ Deced			
Database	Password.	•	47	Test Connectivity @ Passed			
					Cance	Set Whitelist and Next	
Section Parameter		Description					
N/A Task Name		DTS automatica specify an inforr	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to				

specify a unique task name.

Section	Parameter	Description				
		Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database with Public IP Address</b> .				
	Instance Type	<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .				
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to configure the <b>Instance Region</b> parameter.				
	Instance Region	<b>Note</b> If a whitelist is configured for the self-managed PostgreSQL database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click <b>Get IP Address Segment of DTS</b> next to <b>Instance Region</b> to obtain the CIDR blocks of DTS servers.				
Source	Database Type	Select PostgreSQL.				
Database	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed PostgreSQL database. In this example, enter the public IP address.				
	Port Number	Enter the service port number of the self-managed PostgreSQL database. The default port number is <b>5432</b> .				
	Dat abase Name	Enter the name of the self-managed PostgreSQL database.				
	Dat abase Account	Enter the account that is used to log on to the self-managed PostgreSQL database. For information about the permissions that are required for the account, see Permissions required for database accounts.				
		Enter the password of the database account.				
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.				
	Instance Type	Select RDS Instance.				
	Instance Region	Select the region where the destination RDS instance resides.				

Section	Parameter	Description		
	RDS Instance ID	Select the ID of the destination RDS instance.		
Destinatio n Database		Enter the name of the destination database in the RDS instance. The name can be different from the name of the self-managed PostgreSQL database.		
	Database Name	<b>Note</b> Before you configure the data migration task, you must create a database and schema in the destination RDS instance. For more information, see Before you begin.		
	Database Account	Enter the username of the account that is used to log on to the RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
		Enter the password of the database account.		
	Dat abase Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects that you want to migrate.
| 1.Configure So   | arce and Destination 2.Configure Migration Types and 3.Advanced Settings 4.Precheck  |  |  |  |
|--|--|--|--|--|
| * Migratic<br>triggers, Fo<br>Note: do<br>cleans u     | on Types: Schema Migration Full Data Migration Incremental Data Migration Note: Incremental data migration does not support<br>r more information, see Reference<br>not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database<br>p the log too early, the DTS incremental task may fail |  |  |  |
| Available<br>Expand<br>C C ter                         | Selected (To edit an object name or its filter, hover over the object and dick Edit.) Learn more.  |  |  |  |
| Select All   | Remove All   |  |  |  |
| *Rename [<br>* Retry Tin                               | Databases and Tables:     Image Database and Table Names     Change Database and Table Names       The for Failed Connection     720     Minutes   |  |  |  |
| Informati<br>1. Data mig<br>in the sour<br>2. Do not d | on:<br>pration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema<br>te database.<br>to DDL operation during structure and full migration, otherwise the task may fail<br>Cancel Previous Save Precheck   |  |  |  |
| Setting  | Description  |  |  |  |
|  | <ul> <li>To perform only full data migration, select Schema Migration and Full Data<br/>Migration.</li> </ul>  |  |  |  |
| Select<br>the  | <ul> <li>I o ensure service continuity during data migration, select Schema Migration, Full</li> <li>Data Migration, and Incremental Data Migration. In this example, select all of the three migration types.</li> </ul>  |  |  |  |
| n types  | <b>Note</b> If Incremental Data Migration is not selected, do not write data to the self-managed PostgreSQL database during full data migration. This ensures data consistency between the self-managed database and the RDS instance.   |  |  |  |

Setting	Description				
Select the objects that you want to migrate	<ul> <li>Select one or more objects from the Available section and click the ) icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or schemas as the objects to be migrated.</li> <li>After an object is migrated to the destination RDS instance, the name of the object remains the same as that in the self-managed PostgreSQL database. You can use the object name mapping feature to rename the objects that are migrated to the destination RDS instance. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>				
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination RDS instance. For more information, see Object name mapping.				
Specify the retry time range for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS reconnects to the self-managed database or RDS instance within the specified period of time, DTS resumes the data migration task. Otherwise, the data migration task fails.				
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.				

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the



icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the destination RDS instance.

#### What to do next

The database accounts that are used for data migration have the read and write permissions. After data migration is completed, you must delete the accounts of both the self-managed PostgreSQL database and the ApsaraDB RDS for PostgreSQL instance to ensure security.

# 6.4.3. Migrate full data from a self-managed PostgreSQL database to an ApsaraDB RDS for

# PostgreSQL instance

This topic describes how to migrate full data from a self-managed PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. To migrate full data from a self-managed PostgreSQL database, you can select schema migration and full data migration when you configure a task.

# **Background information**

• This topic uses User-Created Database with Public IP Address as an example to describe how to

configure a full data migration task. To ensure data consistency, we recommend that you do not write data to the self-managed PostgreSQL database during full data migration. For information about how to migrate data with minimal downtime, see Migrate incremental data from a self-managed PostgreSQL database (version 10.1 to 13) to an ApsaraDB RDS for PostgreSQL instance and Migrate incremental data from a self-managed PostgreSQL database (in PostgreSQL 10.0 or an earlier version) to an ApsaraDB RDS for PostgreSQL instance.

• After you migrate full data from a self-managed PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance, you can restore data by using logical backup files. For more information, see Restore data from a logical backup file.

# Prerequisites

- The version of the self-managed PostgreSQL database is 9.2, 9.3, 9.4, 9.5, 9.6, 10.x, 11, 12, or 13.
- The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total size of the data in the self-managed PostgreSQL database.
- The service port of the self-managed PostgreSQL database is accessible over the Internet.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The name of the source database cannot contain hyphens (-), for example, dts-testdata.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- A single data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each database.
- To ensure that the data migration task runs as expected, you can perform a primary/secondary switchover only on a V11 ApsaraDB RDS for PostgreSQL instance. In this case, you must set the rds\_failover\_slot\_mode parameter to sync . For more information, see Logical Replication Slot Failover.

Q Warning If you perform a primary/secondary switchover on a self-managed PostgreSQL database or an ApsaraDB RDS for PostgreSQL instance of a version other than V11, the data migration task stops.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.

# **Migration types**

• Schema migration

DTS migrates the schemas of the required objects to the destination database. DTS supports schema migration for the following types of objects: table, trigger, view, sequence, function, user-defined type, rule, domain, operation, and aggregate.

#### • Full data migration

DTS migrates historical data of the required objects from the self-managed PostgreSQL database to the destination database in the ApsaraDB RDS for PostgreSQL instance.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Pricing.

## Permissions required for database accounts

Database	Schema migration	Full data migration
Self-managed PostgreSQL database	The USAGE permission on pg_catalog	The SELECT permission on the objects to be migrated
ApsaraDB RDS for PostgreSQL instance	The CREATE and USAGE permissions on the objects to be migrated	The permissions of the schema owner

For information about how to create and authorize a database account, see the following topics:

- Self-managed PostgreSQL database: CREATE USER and GRANT
- ApsaraDB RDS for PostgreSQL instance: Create an account on an ApsaraDB RDS for PPAS instance.

# Process of full data migration

To prevent data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

1. Migrate the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates.

**?** Note Functions that are written in the C programming language cannot be migrated.

- 2. Migrate full data.
- 3. Migrate the schemas of triggers and foreign keys.

## Before you begin

Create a database and schema in the destination RDS instance based on the database and schema information of the objects to be migrated. The schema name of the source and destination databases must be the same. For more information, see Create a database on an ApsaraDB RDS for PostgreSQL instance and Appendix: User and schema management.

## Procedure

1. Loa on to the DTS console.

- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and De	stination [	Databases 2.Con	figure Migration Types and Objects		3.Map name modification	A.Precheck
* Task	Name:	PostgreSQL_TO_RDS				
Source Database						
* Instan	ice Type:	User-Created Database w	ith Public IP Address	*		
* Instance	e Region:	Singapore		•	Get IP Address Segment of DTS	
* Databa	ase Type:	PostgreSQL		٣		
* Hostname or IP	Address:				]	
* Port	Number:	5432				
* Databas	se Name:	mytestdata			]	
* Database	Account:	dtstest			]	
* Database P	assword:	********		\$	Test Connectivity 🔗 Passed	
Destination Database						
* Insta	nce Type:	RDS Instance		Ŧ		
* Instanc	e Region:	Singapore		•		
* RDS Ins	stance ID:			•		
* Databa	ase Name:	mytestdata		_		
* Database	e Account:	dtstest				
* Database I	Password:	*****	(	<b>\$</b> >	Test Connectivity 📀 Passed	
						Cancel Set Whitelist and Next
Section	Pai	rameter	Description			
N/A	Ta	sk Name	DTS automatica specify an inforr use a unique tas	lly g nativ k na	enerates a task name. W ve name to identify the t me.	'e recommend that you ask. You do not need to
			Select an instand database. In this <b>Public IP Addre</b>	ce ty s exa ess.	rpe based on the deploy Imple, select <b>User-Crea</b> t	ment of the source ted Database with
	Instance Type		<b>? Note</b> If the network e information, s	you nvira ee P	a select other instance ty onment for the source da reparation overview.	rpes, you must deploy atabase. For more

Section	Parameter	Description			
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to specify the <b>Instance Region</b> parameter.			
Source Dat abase	Instance Region	<b>? Note</b> If a whitelist is configured for the self-managed PostgreSQL database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click <b>Get IP</b> <b>Address Segment of DTS</b> next to <b>Instance Region</b> to obtain the CIDR blocks of DTS servers.			
	Database Type	Select PostgreSQL.			
	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed PostgreSQL database. In this example, enter the public IP address.			
	Port Number Enter the service port number of the self-managed PostgreSQL database. The default port number is <b>5432</b> .				
	Dat abase Name	Enter the name of the self-managed PostgreSQL database.			
	Database Account	Enter the account that is used to log on to the self-managed PostgreSQL database. For information about the permissions that are required for the account, see Permissions required for database accounts.			
		Enter the password of the database account.			
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			
	Instance Type	Select RDS Instance.			
	Instance Region	Select the region where the destination RDS instance resides.			
	RDS Instance ID	Select the ID of the destination RDS instance.			

Section	Parameter	Description			
Destinatio n		Enter the name of the destination database in the RDS instance. The name can be different from the name of the self-managed PostgreSQL database.			
Database	Dat abase Name	<b>Note</b> Before you configure the data migration task, create a database and schema in the destination RDS instance. For more information, see Before you begin.			
	Dat abase Account	Enter the database account of the destination RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.			
		Enter the password of the database account.			
	Dat abase Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

**Note** DTS adds the CIDR blocks of DTS servers to the whitelist of the destination RDS instance. This ensures that DTS servers can connect to the destination RDS instance.

#### 7. Select the migration types and the objects to be migrated.

1.Configure So	urce and Destination 2.Configure Migration Types and 3.Advanced Settings 4.Precheck			
<ul> <li>Migration</li> <li>triggers. For</li> <li>Note: do</li> <li>cleans un</li> </ul>	on Types: 🗹 Schema Migration 🔽 Full Data Migration 🔽 Incremental Data Migration Note: Incremental data migration does not support Ir more information, see Reference In not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database In the log too early, the DTS incremental task may fail			
Available Expand • • • • •	Selected (To edit an object name or its filter, hover over the object and dick   the tree before you perform a gloi   Q   stdb     Image: Control of the image o			
Select All	Remove All			
*Rename Databases and Tables:				
	Cancel Previous Save Precheck			
Setting	Description			
Select he nigratio	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> <li>In this example, you must select Schema Migration and Full Data Migration.</li> </ul>			
ı types	<b>Note</b> To ensure data consistency, we recommend that you do not write data to the self-managed PostgreSQL database during full data migration			

Setting	Description				
Select the objects to be migrate d	Select one or more objects from the <b>Available</b> section and click the > icon to add the objects to the <b>Selected</b> section.				
	<ul> <li>Note</li> <li>You can select columns, tables, or schemas as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>				
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.				
Specify the retry time for failed connecti ons to the source or destinati on	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.				
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.				
databas e					

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.

11. Click **Buy and Start** to start the data migration task.

Note We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.

12. Switch your workloads to the destination RDS instance.

# What to do next

The database accounts that are used for data migration have the read and write permissions. After data migration is completed, you must delete the accounts of both the self-managed PostgreSQL database and the ApsaraDB RDS for PostgreSQL instance to ensure security.

# 6.5. Source database: Redis

# 6.5.1. Migrate data from a self-managed Redis

# database to an ApsaraDB for Redis instance

This topic describes how to migrate data from a self-managed Redis database to an ApsaraDB for Redis instance by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration. When you migrate data from a self-managed Redis database to Alibaba Cloud, you can select the two migration types to ensure service continuity.

# Prerequisites

- The version of the self-managed Redis database is 2.8, 3.0, 3.2, 4.0, 5.0, or 6.0.
- The self-managed Redis database is deployed in the standalone architecture rather than the cluster architecture.

**Note** If the self-managed Redis database is deployed in the cluster architecture, you can migrate data by using the data synchronization feature. For more information, see Synchronize data from a self-managed Redis cluster to an ApsaraDB for Redis cluster instance.

- The PSYNC or SYNC command can be run on the self-managed Redis database.
- The available storage space of the destination ApsaraDB for Redis database is larger than the total size of the data in the self-managed Redis database.

# Precautions

• DTS consumes the resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If you migrate a large volume of data or if the server specifications do not meet your requirements, database services may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.

•

• If an expiration policy is enabled for some keys in the source database, these keys may not be deleted in a timely manner after they expired. Therefore, the number of keys in the destination database may be less than that in the source database. You can run the info command to view the number of keys in the destination database.

**?** Note The number of keys that do not have an expiration policy or have not expired is the same in the source and destination databases.

- If you run the EVAL or EVALSHA command to call Lua scripts, DTS cannot identify whether these Lua scripts are executed on the destination database. During incremental data migration, the destination database does not explicitly return the execution results of Lua scripts.
- When you run the PSYNC or SYNC command to transfer data of the LIST type, DTS does not perform the flush operation on the existing data. Therefore, the destination database may contain duplicate data records.
- During data migration, if the number of shards in the self-managed Redis database is increased or decreased, or if the specifications of the database are changed (for example, the memory capacity is scaled up), you must reconfigure the task. To ensure data consistency, we recommend that you clear the data that has been migrated to the destination Redis database before you reconfigure the task.
- During data migration, if the endpoint of the self-managed Redis database is changed, you must submit a ticket to update the change. Otherwise, the append-only files (AOF) of the self-managed Redis database may be reset. In this case, you must reconfigure the task.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.
- Limits on migrating a standalone ApsaraDB for Redis instance to an ApsaraDB for Redis cluster instance: Each command can be run only on a single slot in an ApsaraDB for Redis cluster instance. If you perform operations on multiple keys in the source database and the keys belong to different slots, an error occurs.

CROSSSLOT Keys in request don't hash to the same slot

We recommend that you perform operations on only one key during data migration. Otherwise, the migration task will be interrupted.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# Migration types

• Full data migration

DTS migrates historical data of the required objects from the self-managed Redis database to the destination ApsaraDB for Redis instance.

**?** Note If you perform only full data migration, we recommend that you do not write data to the self-managed Redis database during full data migration. This ensures data consistency.

• Incremental data migration

After full data migration is completed, DTS synchronizes incremental data from the self-managed Redis database to the destination ApsaraDB for Redis instance. Incremental data migration allows you to ensure service continuity when you migrate data from a self-managed Redis database to Alibaba Cloud.

#### Commands that can be migrated

- APPEND
- BIT OP, BLPOP, BRPOP, and BRPOPLPUSH
- DECR, DECRBY, and DEL
- EVAL, EVALSHA, EXEC, EXPIRE, and EXPIREAT
- FLUSHALL and FLUSHDB
- GEOADD and GET SET
- HDEL, HINCRBY, HINCRBYFLOAT, HMSET, HSET, and HSET NX
- INCR, INCRBY, and INCRBYFLOAT
- LINSERT, LPOP, LPUSH, LPUSHX, LREM, LSET, and LT RIM
- MOVE, MSET, MSET NX, and MULTI
- PERSIST, PEXPIRE, PEXPIREAT, PFADD, PFMERGE, PSETEX, and PUBLISH
- RENAME, RENAMENX, RESTORE, RPOP, RPOPLPUSH, RPUSH, and RPUSHX
- SADD, SDIFFST ORE, SELECT, SET, SET BIT, SET EX, SET NX, SET RANGE, SINTERST ORE, SMOVE, SPOP, SREM, and SUNIONST ORE
- ZADD, ZINCRBY, ZINTERSTORE, ZREM, ZREMRANGEBYLEX, ZUNIONSTORE, ZREMRANGEBYRANK, and ZREMRANGEBYSCORE

#### Preparations before incremental data migration

To ensure that the incremental data migration task runs as expected, we recommend that you remove the limit on the replication output buffer. This topic uses a server that runs on Linux as an example.

⑦ Note If you perform only full data migration, skip the following steps.

1. Use the redis-cli program to connect to the self-managed Redis database.

**Note** You can use the redis-cli program after you install the Redis client. For more information, visit Redis community official website.

redis-cli -h <host> -p <port> -a <password>

#### ? Note

- <host>: the endpoint that is used to connect to the self-managed Redis database. You can use 127.0.0.1 in this example.
- <port>: the service port number of the self-managed Redis database. The default port number is 6379.
- <password>: the password of the self-managed Redis database.

Example:

redis-cli -h 127.0.0.1 -p 6379 -a Test123456

2. Run the following command to remove the limit on the replication output buffer:

config set client-output-buffer-limit 'slave 0 0 0'

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source ar	nd Destinatior	2.Configu	ure Migration Types and Objects >		3.Map name modificatio	on >	4.Precheck
*	Task Name: Re	edis					
Source Database							
				_			
*1	Instance Type:	User-Created Databa	se with Public IP Address 🔹	'			
* In:	stance Region:	Singapore	¥	'	Get IP Address Segment of	f DTS	
* 0	atabase Type:	Redis	Ŧ	,			
* I	nstance Mode:	Standalone		7			
" Hostname	or IP Address:						
	Port Number:	6379				() Proved	
Datab	ase Password:	*****	¢>	<u>'</u>	Test Connectivity	✓ Passed	
Destination Database							
* 1	Instance Tuner	De die Testen en		_			
* To	stance Pogion:	Cineman Cineman Cineman Cineman Cineman Cinema Cine	-				
* Pod	is Instance ID:	Singapore	•				
Datab	is instance iD.		· ·			() Droved	
Datab	dse Passworu:	*****	ψ	' 	Test Connectivity	@ Passed	
							Cancel Set Whitelist and Next
Section	Param	neter	Description				
N/A	Task Name		DTS automatically gene specify an informative r specify a unique task na	era nai am	ntes a task nar me to identify ne.	me. We re the task.	commend that you You do not need to
	Instance Type		Select an instance type managed Oracle datab Database with Public	as as	ased on the d e. In this exam P Address.	eploymen Iple, selec	nt of the self- ct <b>User-Created</b>
			<b>Note</b> If you see the environment that database. For more in	ele : is nf	ct other instar required for t ormation, see	nce types, the self-m <b>Preparati</b>	, you must set up nanaged Oracle <mark>on overview</mark> .

Section	Parameter	Description				
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to configure the <b>Instance Region</b> parameter.				
	Instance Region	<b>?</b> Note If a whitelist is configured for the self-managed Redis database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click <b>Get IP Address</b> <b>Segment of DTS</b> next to <b>Instance Region</b> to obtain the CIDR blocks of DTS servers.				
Source	Database Type	Select Redis.				
Database	Instance Mode	The value of this parameter is set to <b>Standalone</b> and cannot be changed to Cluster.				
	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed Redis database. In this example, enter the public IP address.				
	Port Number	Enter the service port number of the self-managed Redis database. The default port number is 6379. <b>Note</b> The service port of the self-managed Redis database must be accessible over the Internet.				
		Enter the password of the self-managed Redis database.				
	Dat abase Password	<b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.				
	Instance Type	Select Redis Instance.				
	Instance Region	Select the region where the destination ApsaraDB for Redis instance resides.				
	Redis Instance ID	Select the ID of the destination ApsaraDB for Redis instance.				
Destinatio n Database						

Section	Parameter	Description		
Dat abase Password	Enter the database password of the destination ApsaraDB for Redis instance.			
	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### ☐ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects that you want to migrate.

1.Configure Sou	arce and Destination 2.Configure Migration Types and 3.Advanced Settings 4.Precheck				
<ul> <li>Migratio</li> <li>triggers, Fo</li> <li>Note: do</li> <li>cleans up</li> </ul>	n Types: Schema Migration Full Data Migration Incremental Data Migration Note: Incremental data migration does not support rmore information, see Reference not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database the log too early, the DTS incremental task may fail				
Available Expand Expand The test Select All *Rename D * Retry Tim Informatio 1. Data mig in the sour 2. Do not d	selected (To esit an object name or its filter, hover over the object and click     titl23   ittb     ittb				
	Cancel Previous Save Precheck				
Setting	Description				
Select the	<ul> <li>To perform only full data migration, select only Full Data Migration.</li> <li>To ensure service continuity during data migration, select Full Data Migration and Incremental Data Migration.</li> </ul>				
migratio n types	<b>?</b> Note If Incremental Data Migration is not selected, we recommend that you do not write data to the self-managed Redis database during data migration. This ensures data consistency between the source and destination databases.				
Select the objects	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.				
that you want to migrate	<b>Note</b> You can select only databases as the objects to be migrated.				

Setting	Description
Specify whether to rename objects	In this scenario, you cannot rename objects.
Specify the retry time range for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS reconnects to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

#### 8. Click Precheck.

- ? Note
  - A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
  - If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

Task Name   Search by migration task name.	Search Sort: Default Sorting V Status: All V	⊎ Tag
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Ta	sk Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed Incremental Data M 1 n The migration task is not delayed.
Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) <

12. Switch your workloads to the destination ApsaraDB for Redis instance.

## What to do next

The database accounts that are used for data migration have the read and write permissions. After the data migration task is completed, you must change the passwords of the self-managed Redis database and the ApsaraDB for Redis instance to ensure database security.

# 6.6. Source database: MongoDB

# 6.6.1. Migrate a self-managed standalone

# MongoDB database to Alibaba Cloud by using

# DTS

This topic describes how to migrate data from a self-managed standalone MongoDB database to an ApsaraDB for MongoDB instance by using Data Transmission Service (DTS).

You can also use the built-in commands of MongoDB to migrate self-managed MongoDB databases. For more information, see Migrate self-managed MongoDB databases to standalone instances by using tools provided by MongoDB.

For more information about data migration or synchronization solutions, see Overview.

## Prerequisites

- The versions of the source and destination MongoDB databases are supported by DTS. For more information, see Overview of data migration scenarios.
- The available storage space of the ApsaraDB for MongoDB instance is larger than the total size of the data in the self-managed MongoDB database.

#### Precautions

• By default, oplogs are not supported by self-managed standalone MongoDB databases. If you use a replica set that contains only one node, oplogs are supported for DTS to implement incremental

migration.

- To prevent service disruptions, we recommend that you migrate data during off-peak hours.
- If the source and destination MongoDB databases use different versions or storage engines, make sure that your applications can run on both databases. For more information about the versions and storage engines that are supported by ApsaraDB for MongoDB, see MongoDB versions and storage engines.
- The admin or local database cannot be used as the source or destination database.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba
Incremental data migration	Charged. For more information, see Pricing.	Cloud over the Internet. For more information, see Pricing.

# Migration types

• Full data migration: DTS migrates all historical data of the source MongoDB database to the destination MongoDB database.

**?** Note DTS can migrate the following types of objects: database, collection, and index.

• Incremental data migration: After full data migration is completed, DTS synchronizes incremental data of the source MongoDB database to the destination MongoDB database.

#### ? Note

- DTS can synchronize the create and delete operations that are performed on databases, collections, and indexes.
- DTS can synchronize the create, delete, and update operations that are performed on documents.

## Permissions required for database accounts

Database	Full data migration
Self-managed MongoDB database	The read permissions on the source database
ApsaraDB for MongoDB instance	The read and write permissions on the destination database

For more information about how to create and authorize a database account, see the following topics:

- Self-managed MongoDB database: db.createUser()
- ApsaraDB for MongoDB instance: Manage user permissions on MongoDB dat abases

# Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the ApsaraDB for MongoDB instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination D	atabases 2.Configure Migration Types and Objects	3.Advanced Settings	A.Precheck
* Task Name:			
Source Database			
* Instance Type:	User-Created Database with Public IP Address	DTS support type	
* Instance Region:	China (Hangzhou)	Get IP Address Segment of DTS	
* Database Type:	MongoDB v		
Hostname or IP Address:			
* Port Number:	27017		
Database Name:	admin	Authenticate Database with Account	
Database Account:	-		
Database Password:	ه	Test Connectivity 🔗 Passed	
* Encryption:	Non-encrypted OSL-encrypted(MongoDB Atlas only)		
Destination Database			
* Instance Type:	MongoDB Instance v		
* Instance Region:	China (Hangzhou)		
* MongoDB Instance ID:	dds-bp:		
* Database Name:	admin	Authenticate Database with Account	
* Database Account:	root		
* Database Password:	ه>	Test Connectivity 🔗 Passed	
			Cancel Set Whitelist and Next

Section	Parameter	Description	
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to use a unique task name.	
		Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database with Public IP Address</b> .	
Instance Type	<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .		

Section	Parameter	Description
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to specify the <b>Instance Region</b> parameter.
	Instance Region	Note If a whitelist is configured for the self-managed MongoDB database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
	Database Type	Select MongoDB.
Source	Hostname or IP Address	Enter the endpoint that is used to connect to the self-managed MongoDB database. In this example, enter the public IP address.
Dat abase	Port Number	Enter the service port number of the self-managed MongoDB database.          Image: The service port of the self-managed MongoDB database must be accessible over the Internet.
	Database	Enter the name of the authentication database. The database account
	Name	is created in this database.
	Dat abase Account	Enter the account that is used to log on to the self-managed MongoDB database. For information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the database account.
	Dat abase Password	<b>?</b> Note After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.
		Select Non-encrypted.
	Encryption	<b>Note</b> You can select <b>SSL-encrypted</b> only when you migrate data from MongoDB Atlas.
	Instance Type	Select MongoDB Instance.

Section	Parameter	Description		
Destinatio n Database	Instance Region	Select the region where the destination ApsaraDB for MongoDB instance resides.		
	MongoDB Instance ID	Select the ID of the destination ApsaraDB for MongoDB instance.		
	Dat abase Name	Enter the name of the authentication database. The database account is created in this database.		
	Database Account	Enter the database account of the destination ApsaraDB for MongoDB instance. For information about the permissions that are required for the account, see Permissions required for database accounts.		
	Database Password	Enter the password of the database account. <b>Note</b> After you specify the destination database parameters, click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.		

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### ☐ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

Migration Types: Full Data Migration eference Note: do not clean up the incremental data cleans up the log too early, the DTS increm	✓ Incremental Data Migration Incremental Data Migration log generated by the source data entail task may fail	Note: Increm	ental data migration does	not support trigger DTS full task is run	rs. For more information
Available		1	Selected		
Expand the tree before you perform a glo	Q		nongodbtest		
		> <			
Select All	•		Remove All		
Rename Databases and Tables	Do Not Change Database and T	able Names	Channe Datahase and	Table Names	
Retry Time for Failed Connection <b>nformation:</b> . Data migration only copies the data and sch the source database. . Do not do DDL operation during structure a	720 Minutes (7) ema in the source database and s	aves the copy in t	he destination database. T	The process does n	ot affect any data or sch
			Cancel	Previous	Save Precheck

Setting	Description		
	<ul> <li>To perform only full data migration, select only Full Data Migration.</li> <li>To ensure service continuity during data migration, select Full Data Migration and Incremental Data Migration.</li> </ul>		
Select migratio n types	<ul> <li>Note</li> <li>To migrate the incremental data of a standalone MongoDB database by using DTS, you must enable the oplog feature for the database.</li> <li>If Incremental Data Migration is not selected, we recommend that you do not write data to the source instance during full data migration. This ensures data consistency between the source and destination instances.</li> </ul>		
	• Select one or more objects from the <b>Available</b> section and click the <b>S</b> icon to move		
	the objects to the <b>Selected</b> section.		
Select the objects to be migrate d	<b>Note</b> DTS cannot migrate data from the admin, local, or config database.		
	<ul> <li>You can select databases, collections, or functions as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> </ul>		
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.		

Setting	Description
Specify the retry time range for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time range based on your business requirements. If DTS reconnects to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
ons to the source or destinati on databas e	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can troubleshoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the data migration task automatically stops.

• Incremental data migration

The task does not automatically stop during incremental data migration. You must manually stop the task.

**Note** We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the ApsaraDB for MongoDB instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed are displayed in the progress bar of the data migration task. Then, stop writing data to the source instance for a few minutes. The latency of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

Task ID/Name:	Status: Migrating	Pause Task   View Details   Duplicate Task   Norrade   Configure Monitoring and Alerting
2 2019-07-04 15:57:43 Created Full Data Migration 100 3 grated Rows: 5000003)		Completed Incremental Data Migration The migration task is not delayed.
Start Pause Stop Delete		Total: 1 item(s) , Per Page: 20 item(s) ( < 1 > >

12. Switch your workloads to the destination ApsaraDB for MongoDB instance.

# 6.6.2. Migrate a self-managed MongoDB database that uses the replica set architecture to ApsaraDB for MongoDB by using DTS

This topic describes how to use Data Transmission Service (DTS) to migrate a self-managed MongoDB database that uses the replica set architecture to ApsaraDB for MongoDB. DTS supports full data migration and incremental data migration. When you configure a data migration task, you can select these two migration types to ensure service continuity.

You can also use the built-in commands of MongoDB to migrate self-managed MongoDB databases. For more information, see Migrate self-managed databases to Alibaba Cloud by using tools provided by MongoDB. For more information about data migration and synchronization solutions, see Overview.

## Prerequisites

- The versions of the source and destination MongoDB databases are supported by DTS. For more information, see Overview of data migration scenarios.
- The available storage space of the ApsaraDB for MongoDB instance is larger than the total size of the data in the self-managed MongoDB database.

## Precautions

- DTS consumes the resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If you migrate a large amount of data or if the server specifications do not meet your requirements, database services may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- If the source and destination databases have different versions or storage engines, make sure that the versions or storage engines are compatible. For more information, see MongoDB versions and storage engines.
- The admin or local database cannot be used as the source or destination database.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba
Incremental data migration	Charged. For more information, see Pricing.	Cloud over the Internet. For more information, see Pricing.

# **Migration types**

• Full data migration: All existing data in the source instance is migrated to the destination instance.

**?** Note DTS can migrate the following types of objects: database, collection, and index.

• Incremental data migration: After full data migration is complete, DTS synchronizes incremental data of the source MongoDB database to the destination MongoDB database.

? Note

- DTS can synchronize the create and delete operations that are performed on databases, collections, and indexes.
- DTS can synchronize the create, delete, and update operations that are performed on documents.

# Permissions required for database accounts

Database	Full data migration	Incremental data migration
Self-managed MongoDB database	Read permissions on the source database	Read permissions on the source, admin, and local databases
ApsaraDB for MongoDB instance	Read and write permissions on the destination database	Read and write permissions on the destination database

For more information about how to create a database account and grant permissions to the account, see the following topics:

- ApsaraDB for MongoDB instances: Manage user permissions on MongoDB databases
- Self-managed MongoDB database: db.createUser()

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. On the upper part of the **Migration Tasks** page, select the region where the destination ApsaraDB for MongoDB instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination D	atabases 2.Configure Migration Types and Objects	3.Advanced Settings	4.Precheck
* Task Name:			
Source Database			
* Instance Type:	User-Created Database with Public IP Address	DTS support type	
* Instance Region:	China (Hangzhou)	Get IP Address Segment of DTS	
* Database Type:	MongoDB <b>v</b>		
Hostname or IP Address:		]	
* Port Number:	27017		
Database Name:	admin	Authenticate Database with Account	
Database Account:		]	
Database Password:		Test Connectivity	
* Encryption:	Non-encrypted      SSL-encrypted(MongoDB Atlas only)		
Destination Database			
Instance Type:	MonooDB Instance		
* Instance Region:	China (Hangzhou)		
* MongoDB Instance ID:	dds-bp: 🗸		
* Database Name:	admin	Authenticate Database with Account	
* Database Account:	root	]	
* Database Password:	مله (	Test Connectivity	
			Cancel Set Whitelist and Next

Section	Parameter	Description		
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.		
		The access method of the source database. In this example, <b>User-</b> <b>Created Database with Public IP Address</b> is selected.		
	Instance Type	<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .		
		If you select <b>User-Created Database with Public IP Address</b> for the Instance Type parameter, you do not need to configure the <b>Instance Region</b> parameter.		
	Instance Region	Note If a whitelist is configured for the self-managed MongoDB database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.		
	Database Type	Select MongoDB.		

#### Dat a Transmission Service

Section	Parameter	Description		
	Hostname or IP Address	The endpoint that is used to connect to the self-managed MongoDB database. In this example, enter the public IP address.		
Source		The service port number of the self-managed MongoDB database.		
Database	Port Number	<b>Note</b> The service port of the self-managed MongoDB database must be accessible over the Internet.		
	Database Name	The name of the authentication database. The database account is created in this database.		
	Database Account	The account that is used to log on to the self-managed MongoDB database. For information about the permissions that are required for the account, see Permissions required for database accounts.		
		The password of the database account.		
	Dat abase Password	<b>?</b> Note After you specify the information about the source database, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is correct, the <b>Passed</b> message is displayed. If the information is incorrect, the <b>Failed</b> message is displayed and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.		
		Select Non-encrypted.		
	Encryption	<b>Note</b> You can select <b>SSL-encrypted</b> only when you migrate data from MongoDB Atlas.		
	Instance Type	Select MongoDB Instance.		
	Instance Region	The region where the destination ApsaraDB for MongoDB instance is deployed.		
	MongoDB Instance ID	The ID of the destination instance.		
	Database Name	The name of the authentication database. The database account is created in this database.		
	Dat abase Account	The database account of the destination ApsaraDB for MongoDB instance. For information about the permissions that are required for the account, see Permissions required for database accounts.		
Destinatio n Database				

Section	Parameter	Description
		The password of the database account.
	Database Password	Note After you specify the information about the destination instance, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is correct, the <b>Passed</b> message is displayed. If the information is incorrect, the <b>Failed</b> message is displayed and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### ☐ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and objects to migrate.

1.Configu	re Source and 🔰 2	Configure Migration Types and		3.Advanced Settings	$\rightarrow$	4.Precheck
<ul> <li>Migration Typ</li> <li>Reference</li> </ul>	es: 🗸 Full Data Migration	✓ Incremental Data Migration	Note: Incre	emental data migration does	not support trigg	ers. For more information, see
Note: do not o cleans up the	lean up the incremental data log too early, the DTS increm	log generated by the source databa ental task may fail	se after the D	TS task is started when the	DTS full task is ru	nning. If the source database
Available				Selected		
Expand the tr	ee before you perform a glol	Q				Q
📔 admin				e mongodbtest		
			>			
			`			
		-				
Select All				Remove All		
Rename Databa	ises and Tables:	Do Not Change Database and Tab	ole Names	Change Database and	Table Names	
information:	Palled Connection	720 Minutes				
1. Data migration in the source dat 2. Do not do DDI	only copies the data and sch abase. . operation during structure a	ema in the source database and sain and full migration, otherwise the task	ves the copy ir : may fail	n the destination database. '	The process does	not affect any data or schema
				Cancel	Previous	Save Precheck
Setting	Description					
	• To perform o	only full data migratic	on, select	t only <b>Full Data</b>	Migration	
	<ul> <li>To ensure se</li> <li>Incremental</li> </ul>	rvice continuity during l <b>Data Migration</b> .	) data m	igration, select F	ull Data M	ligration and

Note If Incremental Data Migration is not selected, we recommend that you do not write data to the self-managed MongoDB database during full data migration. This ensures data consistency between the source and destination databases.

Select

migratio

n types

the

Setting	Description
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to add the objects to the Selected section.</li> <li>Note DTS cannot migrate data from the admin, local, or config database.</li> <li>The objects that you can migrate are databases, collections, and functions.</li> <li>By default, after an object is migrated to the destination instance, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
ons to the source or destinati on	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
databas e	

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Instance Class** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.

- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

The task does not automatically stop during incremental data migration. You must manually stop the task.

**?** Note We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the ApsaraDB for MongoDB instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed are displayed in the progress bar of the data migration task. Then, stop writing data to the source database for a few minutes. The latency of incremental data migration may be displayed in the progress bar.
- b. Wait until the state of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

Task ID/Name:	Status: Migrating	Pause Task   View Details   Duplicate Task   Hograde   Configure Monitoring and Alerting
2 2019-07-04 15:57:43 Created Full Data Migration 100 3 grated Rows: 5000003)		Completed Incremental Data Migration The migration task is not delayed.
Start Pause Stop Delete		Total: 1 item(s) , Per Page: 20 item(s) $\begin{tabular}{ c c c c } \hline & & & & \\ \hline & & & & \\ \hline & & & & \\ \hline & & & &$

12. Switch your workloads to the destination ApsaraDB for MongoDB instance.

## References

Connect to a replica set instance

# 6.6.3. Migrate a self-managed MongoDB database that uses the sharded cluster architecture to ApsaraDB for MongoDB by using DTS

This topic describes how to use Data Transmission Service (DTS) to migrate a self-managed MongoDB database that uses the sharded cluster architecture to an ApsaraDB for MongoDB sharded cluster instance. DTS allows you to migrate the historical and incremental data of on-premises databases to Alibaba Cloud without service disruptions.

For more information about data migration and synchronization solutions, see Overview.

## Prerequisites

- The versions of the source and destination MongoDB databases are supported by DTS. For more information, see Overview of data migration scenarios.
- Each shard in the destination sharded cluster instance has sufficient storage space.

**Note** For example, a self-managed MongoDB database has three shards, and one of these shards occupies a maximum storage space of 500 GB. In this case, the storage space of each shard in destination instance must be larger than 500 GB.

• The admin or local database is not used as the source or destination database.

#### How it works

DTS migrates a self-managed MongoDB database by migrating each shard in the database. You must create a data migration task for each shard.

**Note** The distribution of migrated data in the destination ApsaraDB for MongoDB instance is based on the shard key that you specify. For more information, see **Configure sharding to maximize the performance of shards**.



## Precautions

- DTS consumes the resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If you migrate a large volume of data or the server specifications cannot meet your requirements, database services may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- If the source and destination ApsaraDB for MongoDB instances have different versions or storage engines, make sure that the versions or storage engines are compatible. For more information, see MongoDB versions and storage engines.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	
Incremental data migration	Charged. For more information, see Data Transmission Service Pricing.	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Data Transmission Service Pricing.

# **Migration types**

• Full data migration: DTS migrates all historical data of the source MongoDB database to the destination MongoDB database.

**?** Note DTS can migrate the following types of objects: database, collection, and index.

• Incremental data migration: After full data migration is complete, DTS synchronizes incremental data of the source MongoDB database to the destination MongoDB database.

#### ⑦ Note

- DTS can synchronize the create and delete operations that are performed on databases, collections, and indexes.
- DTS can synchronize the create, delete, and update operations that are performed on documents.

# Permissions required for database accounts

Database	Full data migration	Incremental data migration
Self-managed MongoDB database	The read permissions on the source database	The read permissions on the source database, the admin database, and the local database
ApsaraDB for MongoDB instance	The read and write permissions on the destination database	The read and write permissions on the destination database

For information about how to create and authorize a database account, see the following topics:

- Self-managed MongoDB database: db.createUser()
- ApsaraDB for MongoDB instance: Manage user permissions on MongoDB databases

## Before you begin

1. Disable the balancer of the self-managed MongoDB database. This prevents the impact of chunk migration on data consistency. For more information, see Manage the ApsaraDB for MongoDB balancer.

• Warning If the balancer is not disabled, chunk migration affects the consistency of the data read by DTS.

2. Delete the orphaned documents that are generated due to chunk migration failures from the self-
#### managed MongoDB database.

(?) Note If you do not delete the orphaned documents, the migration performance will be compromised. In addition, some documents may have duplicate \_\_id values and unwanted data may be migrated.

i. Download the cleanupOrphaned.js file.

wget "http://docs-aliyun.cn-hangzhou.oss.aliyun-inc.com/assets/attach/120562/cn\_zh/ 1564451237979/cleanupOrphaned.js"

ii. Replace *test* in the cleanupOrphaned.js file with the name of the database from which you want to delete orphaned documents.

Onte If you want to delete orphaned documents from multiple databases, repeat Steps ii and iii.

```
function cleanupOrphaned(coll) {
  var nextKey = { };
  var result;

while ( nextKey != null ) {
   result = db.adminCommand( { cleanupOrphaned: coll, startingFromKey: nextKey } );

   if (result.ok != 1)
      print("Unable to complete at this time: failure or timeout.")
      printjson(result);
      nextKey = result.stoppedAtKey;
   }
}
var dbName = 'test'
db = db.getSiblingDB(dbName)
db.getCollectionNames().forEach(function(collName) {
            cleanupOrphaned(dbName + "." + collName);
});
```

iii. Run the following command on a shard to delete the orphaned documents from all collections in the specified database:

⑦ Note You must repeat this step on each shard.

mongo --host <Shardhost> --port <Primaryport> --authenticationDatabase <database>
-u <username> -p <password> cleanupOrphaned.js

#### ? Note

- Shardhost>: the IP address of the shard.
- <Primaryport>: the service port of the primary node in the shard.
- <database>: the name of the authentication database. The database account is created in this database.
- <username>: the account that is used to log on to the self-managed MongoDB database.
- <password>: the password that is used to log on to the self-managed MongoDB database.

#### Example:

In this example, a self-managed MongoDB database has three shards, and you must delete the orphaned documents on each shard.

```
mongo --host 172.16.1.10 --port 27018 --authenticationDatabase admin -u root -p 'T
est123456' cleanupOrphaned.js
```

```
mongo --host 172.16.1.11 --port 27021 --authenticationDatabase admin -u root -p 'Te
st123456' cleanupOrphaned.js
```

```
mongo --host 172.16.1.12 --port 27024 --authenticationDatabase admin -u root -p 'T
est123456' cleanupOrphaned.js
```

3. Create databases and collections to be sharded in the destination ApsaraDB for MongoDB instance, and configure data sharding based on your business requirements. For more information, see Configure sharding to maximize the performance of shards.

(?) Note If you configure data sharding before you start data migration, data in the selfmanaged MongoDB database is evenly migrated to the shards in the destination sharded cluster instance. This prevents the overloading of a single shard.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the destination ApsaraDB for MongoDB instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination D	atabases 2.Configure Migration Types and Objects	3.Advanced Settings	4.Precheck
* Task Name:			
Source Database			
* Instance Type:	User-Created Database with Public IP Address	DTS support type	
* Instance Region:	China (Hangzhou)	Get IP Address Segment of DTS	
* Database Type:	MongoDB v		
* Hostname or IP Address:			
* Port Number:	27017		
Database Name:	admin	Authenticate Database with Account	
Database Account:			
Database Password:	····· ﴿>	Test Connectivity Sessed	
* Encryption:	Non-encrypted      SSL-encrypted(MongoDB Atlas only)		
Destination Database			
* Instance Type:	MongoDB Instance 🔻		
* Instance Region:	China (Hangzhou)		
* MongoDB Instance ID:	dds-bp:		
* Database Name:	admin	Authenticate Database with Account	
* Database Account:	root		
* Database Password:	······ 4>	Test Connectivity 🔗 Passed	
			Cancel Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to use a unique task name.
Instance Type		Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database with Public IP Address</b> .
		<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .
		If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to configure the <b>Instance Region</b> parameter.
	Instance Region	<b>Note</b> If a whitelist is configured for the self-managed MongoDB database, you must add the CIDR blocks of DTS servers to the whitelist. You can click <b>Get IP Address Segment of DTS</b> next to <b>Instance Region</b> to obtain the CIDR blocks of DTS servers.
	Database Type	Select MongoDB.

Section	Parameter	Description	
Source Database		Enter the endpoint or IP address of a shard in the self-managed MongoDB database. In this example, enter the public IP address of the shard.	
	Hostname or IP Address	<b>Note</b> DTS migrates each shard of the source database until the whole cluster is migrated. In this example, enter the endpoint or IP address of the first shard. When you configure the second migration task, enter the endpoint or IP address of the second shard. You must repeat this procedure until all shards are migrated.	
		Enter the service port number of the shard.	
	Port Number	<b>Note</b> The service port of each shard in the self-managed MongoDB database must be accessible over the Internet.	
	Database Name	Enter the name of the authentication database. The database account is created in this database.	
	Dat abase Account	Enter the account that is used to log on to the self-managed MongoDB database. For information about the permissions that are required for the account, see Permissions required for database accounts.	
		Enter the password of the database account.	
	Dat abase Password	<b>?</b> Note After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.	
		Select Non-encrypted.	
	Encryption	<b>Note</b> You can select <b>SSL-encrypted</b> only when you migrate data from MongoDB Atlas.	
	Instance Type	Select MongoDB Instance.	
	Instance Region	Select the region where the destination ApsaraDB for MongoDB instance resides.	
	MongoDB Instance ID	Select the ID of the destination sharded cluster instance.	

Section	Parameter	Description	
Destinatio n	Dat abase Name	Enter the name of the authentication database. The database account is created in this database.	
Database	Database Account	Enter the database account of the destination ApsaraDB for MongoDB instance. For information about the permissions that are required for the account, see Permissions required for database accounts.	
	Dat abase Password	Enter the password of the database account.	
		Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.	
		the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.	

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### Q Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.

#### 7. Select the migration types and the objects to be migrated.

1.Configu	re Source and 2.Configure Migration Types and	3.Advanced Settings A.Precheck
<ul> <li>Migration Ty Reference</li> <li>Note: do not cleans up the</li> </ul>	pes: Full Data Migration Incremental Data Migration Note: 1 clean up the incremental data log generated by the source database after the log too early, the DTS incremental task may fail	incremental data migration does not support triggers. For more information, see ne DTS task is started when the DTS full task is running. If the source database
Available		Selected
Expand the t	ree before you perform a glol 🛛 🔍	Q
📔 admin		nongodbtest
	>	
	•	
Select All		Remove All
*Rename Datab * Retry Time for In <b>formation:</b> 1. Data migratio 1. Data migratio 1. Do not do DD	ases and Tables:      O Not Change Database and Table Names      Failed Connection     720     Minutes     1	<ul> <li>Change Database and Table Names</li> <li>py in the destination database. The process does not affect any data or schem</li> </ul>
		Cancel Previous Save Precheck
Setting	Description	
	<ul> <li>To perform only full data migration, sel</li> <li>To ensure service continuity during data Incremental Data Migration.</li> </ul>	ect only <b>Full Data Migration</b> . migration, select <b>Full Data Migration</b> and

Note If Incremental Data Migration is not selected, we recommend that you do not write data to the self-managed MongoDB database during full data migration. This ensures data consistency between the source and destination databases.

Select

migratio

n types

the

Setting	Description
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to add the objects to the Selected section.</li> <li>Note DTS cannot migrate data from the admin, local, or config database.</li> <li>The objects that you can migrate are databases, collections, and functions.</li> <li>By default, after an object is migrated to the destination instance, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti ons to the source or destinati on	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
databas e	

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

1

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-as-you-go) Service Terms**.

- 11. Click **Buy and Start** to start the data migration task.
- 12. Repeat Steps 1 to 11 to create data migration tasks for the remaining shards.
- 13. Stop the data migration tasks.
  - Full data migration

We recommend that you do not manually stop a task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.

• Incremental data migration

A task does not automatically stop during incremental data migration. You must manually stop the task.

**?** Note We recommend that you select an appropriate time to manually stop a data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the ApsaraDB for MongoDB instance.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of all migration tasks. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. After the status of **incremental data migration** changes to **The migration task is not delayed** again, manually stop the migration tasks for all shards.



14. Switch your workloads to the destination ApsaraDB for MongoDB instance.

## 6.7. Source database: TiDB

## 6.7.1. Migrate incremental data from a selfmanaged TiDB database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate incremental data from a self-managed TiDB database to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). Incremental data migration allows you to ensure service continuity when you migrate data to Alibaba Cloud. In this example, Pump, Drainer, and a Kafka cluster are deployed.

## Prerequisites

**Note** Before you migrate incremental data, you can migrate historical data from the selfmanaged TiDB database to the ApsaraDB RDS for MySQL instance. For more information, see Migrate full data from a self-managed TiDB database to an ApsaraDB RDS for MySQL instance.

The destination ApsaraDB RDS for MySQL instance is created. For more information, see Create an ApsaraDB RDS for MySQL instance.

#### ♥ Notice

- The destination ApsaraDB RDS for MySQL instance must reside in the China (Hangzhou), China (Shanghai), China (Qingdao), China (Beijing), China (Shenzhen), China (Zhangjiakou), China (Hong Kong), Singapore (Singapore), US (Silicon Valley), or US (Virginia) region.
- The available storage space of the destination ApsaraDB RDS for MySQL instance must be larger than the total size of the data in the self-managed TiDB database.

## **Background information**



The binary log format and implementation mechanism of a TiDB database are different from those of a MySQL database. To migrate incremental data and minimize modifications to the source TiDB database, you must deploy Pump, Drainer, and a Kafka cluster.

Pump records the binary log files that are generated in TiDB in real time, and sends the binary log files to Drainer. Drainer writes the binary log files to the downstream Kafka cluster. During incremental data migration, DTS retrieves data from the Kafka cluster and migrates the data to the destination database in real time. For example, DTS can migrate incremental data to an ApsaraDB RDS for MySQL instance.

#### Precautions

• DTS uses read and write resources of the source and destination databases during full data

migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.

- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

**(?)** Note For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database on an ApsaraDB RDS for MySQL instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## **Migration types**

Migration type	Description	
	DTS migrates the schemas of required objects to the destination database. DTS supports schema migration for views, tables, and databases.	
Schema migration	A warning TiDB and MySQL are heterogeneous databases. DTS does not ensure that the schemas of the source and destination databases are consistent after schema migration. We recommend that you evaluate the impact of data type conversion on your business. For more information, see Data type mappings between heterogeneous databases.	

Migration type	Description	
Full data migration	DTS migrates the historical data of required objects to the destination database.	
	<b>Note</b> During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.	
Incremental data migration	<ul> <li>DTS retrieves binary log files that are generated in TiDB from the Kafka cluster, and migrates incremental data to the destination database in real time. During incremental data migration, the following SQL operations can be synchronized:</li> <li>DML operations: INSERT, UPDATE, and DELETE</li> <li>DDL operations: CREATE TABLE, DROP TABLE, ALTER TABLE, RENAME TABLE, TRUNCATE TABLE, CREATE VIEW, DROP VIEW, and ALTER VIEW</li> <li>Incremental data migration allows you to ensure service continuity when you migrate data from a self-managed TiDB database to Alibaba Cloud.</li> </ul>	

## Preparations

**?** Note The server on which the source database is deployed must be in the same internal network as the servers on which Pump, Drainer, and the Kafka cluster are deployed. This minimizes the impact of network latency on the incremental data migration task.

- 1. Deploy Pump and Drainer. For more information, see TiDB Binlog Cluster Deployment.
- 2. Modify the configuration file of Drainer and specify a Kafka cluster to receive data from Drainer. For more information, see Binlog Slave Client User Guide.
- 3. Deploy a Kafka cluster by using one of the following methods:
  - Deploy a self-managed Kafka cluster. For more information, visit the Apache Kafka official website.

Warning We recommend that you set the message.max.bytes and replica.fetch. max.bytes parameters for the Kafka broker and the fetch.message.max.bytes parameter for the Kafka consumer to greater values. This ensures that the Kafka cluster can receive the binary log files that are generated in TiDB. For more information, see Kafka 2.5 Documentation.

• Purchase and deploy a Message Queue for Apache Kafka instance. For more information, see Quick start of Message Queue for Apache Kafka.

**?** Note The Message Queue for Apache Kafka instance must be deployed in the same virtual private cloud (VPC) as the source database server. This ensures reliable data transmission and minimizes the impact of network latency on incremental data migration.

- 4. Create a topic in the self-managed Kafka cluster or the Message Queue for Apache Kafka instance.
- 5. Add the CIDR blocks of DTS servers to a whitelist of the TiDB database. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.
  - i. Configure the task name and source database.

* Task Name: T	IDB_To_RDS	
Note richer functions. Click to buy data synchroniza	e: If you need to carry out incremental data migration for a long time, it is n ation Click to see the difference between data migration and synchronization	acommended to use the data synchronization function, which has better network stability and n
Source Database		
* Instance Type:	User-Created Database in ECS Instance 🗸	DTS support type
* Instance Region:	China (Hangzhou)	
* ECS Instance ID:	i-bp 👻	
* Database Type:	TIDB	
* Port Number:	4000	
* Database Account:	dtstest	
* Database Password:	····· Ø	Test Connectivity
* Incremental migration or not:	● Yes ○ No	
* Kafka ClusterType:	User-Created Database in ECS Instance 🗸	
* Instance Region:	China (Hangzhou)	
* ECS Instance ID:	i-bp 🗸	
* KafkaPort Number:	9092	
Kafka Cluster Account:		Optional
Kafka Cluster Password	4>	Optional
* Topic:	tidb_test_topic 🗸	Get Toplic list
	Click Get Topic List and then select the specific topic.	
* Kafka Version	1.0 ~	
* Kafka ClusterEncryption:	Non-encrypted O SCRAM-SHA-256	

Parameter	Description	
Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.	
Instance Type	The access method of the source database. In this example, <b>User-Created Database in ECS Instance</b> is selected.	
	<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see Preparation overview.	
Instance Region	The region of the Elastic Compute Service (ECS) instance on which the source TiDB database is deployed.	
Database Type	Select TiDB.	

Parameter	Description
Port Number	The service port number of the source TiDB database. Default value: 4000.
Database Account	The account of the source TiDB database. The account must have the SELECT permission on the objects to migrate and the SHOW VIEW permission.
Database Password	The password of the database account. Notice After you specify the information about the self-managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
Incremental migration or not	Specifies whether to perform incremental data migration. In this example, Yes is selected. For more information about how to perform only full data migration, see Migrate full data from a self-managed TiDB database to an ApsaraDB RDS for MySQL instance.
Kafka Cluster Type	The access method of the Kafka cluster. In this example, User-Created Database in ECS Instance is selected. If the Kafka cluster is connected over other methods, you must deploy the network environment for the Kafka cluster. For more information, see Preparation overview. Note You cannot select Message Queue for Apache Kafka for the Kafka Cluster Type parameter. If you deploy a Message Queue for Apache Kafka instance, you must select User-Created Database Connected over Express Connect, VPN Gateway, or Smart Access Gateway. Then, you must select the VPC to which the Message Queue for Apache Kafka instance belongs.
Instance Region	The value of this parameter is the same as the region of the source database and cannot be changed.
ECS Instance ID	The ID of the ECS instance that hosts the self-managed Kafka cluster.
Kafka Port Number	The service port number of the self-managed Kafka cluster. Default value: 9092.
Kafka Cluster Account	The username that is used to log on to the Kafka cluster. If no authentication is enabled for the Kafka cluster, you do not need to enter the username.
Kafka Cluster Password	The password that corresponds to the username. If no authentication is enabled for the Kafka cluster, you do not need to enter the password.
Торіс	Click <b>Get Topic List</b> and select a topic name from the drop-down list.
Kafka version	The version of the self-managed Kafka cluster.

Parameter	Description
Kafka Cluster Encryption	Select <b>Non-encrypted</b> or <b>SCRAM-SHA-256</b> based on your business and security requirements.

#### ii. Configure the destination database.

Destination Database					
* Instance Type:	DDC Textenses				
instance type.	RDS Instance	•			
* Instance Region:	China (Hangzhou)	~			
* RDS Instance ID:	rm	•			
* Database Account:	dtstest				
* Database Password:	•••••	<b>(</b> )	Test Connectivity		
* Encryption:	Non-encrypted OSL-encrypted				
				Cancel	Set Whitelist and Next

Parameter	Description			
Instance Type	Select <b>RDS Instance</b> .			
Instance Region	The region where the destination ApsaraDB RDS for MySQL instance resides.			
Database Account	The database account of the destination ApsaraDB RDS for MySQL instance. The account must have read and write permissions on the destination database. For more information about how to create and authorize a database account, see Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account for an ApsaraDB RDS for MySQL instance.			
	The password of the database account.			
Database Password	Notice After you specify the information about the self-managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			
Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> based on your needs. If you select <b>SSL-encrypted</b> , you must enable SSL encryption for the ApsaraDB RDS instance before you configure the data migration task. For more information, see <b>Configure SSL encryption for an ApsaraDB RDS for MySQL instance</b> .			
	<b>Notice</b> The <b>Encryption</b> parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.			

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### □ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and objects to migrate.

Cloud

2.Configure Migration Types and \* Migration Types: 🗹 Schema Migration 🛛 🗹 Full Data Migration 📝 Incremental Data Migration Note: Incremental data migration does not support triggers. For more information, see Reference Note: do not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database cleans up the log too early, the DTS incremental task may fail Available Selected (To edit an object name or its filter, hover over the object and click Edit.) Learn more. Q Expand the tree before you perform a glo Q 🗄 🦢 test123 📔 dtstestdata 🗉 😑 testdb > < Select All Remove All \*Rename Databases and Tables: O Not Change Database and Table Names O Change Database and Table Names \* Retry Time for Failed Connection 720 Minutes 🕐 Information: 1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database. 2. Do not do DDL operation during structure and full migration, otherwise the task may fail ----Cancel Previous Save Prechee

Setting	Description
Select the migratio n types	<ul> <li>If you want to perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>If you want to ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration. In this example, all of the three migration types are selected.</li> </ul>

# Data Migration Migrate data from a self-managed database to Alibaba Cloud

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the > icon to add the objects to the <b>Selected</b> section.
Select the objects that you want to migrate	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated. If you select tables or columns as the objects to be migrated, DTS does not migrate other objects such as views, triggers, and stored procedures to the destination database.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the ApsaraDB RDS instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti ons to the source or destinati on	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
databas e	

## 8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click **Next**.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.

Task Name	✓ Search by migration task name. Se	arch Sort: Default Sorting v Status: All v 🗣 T	ag
Task ID/N	Name: St	atus: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
4 Nov 20 Schema	021.09:59:01 Created a Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed incremental Data M 1 n The migration task is not delayed.
☑ 3	Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) < < 1 > >

## 6.7.2. Migrate full data from a self-managed TiDB

## database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate full data from a self-managed TiDB database to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS).

## Prerequisites

<sup>&</sup>gt; Document Version: 20220712

#### An Create an ApsaraDB RDS for MySQL instance instance is created.

#### ♥ Notice

- The destination ApsaraDB RDS for MySQL instance resides in the China (Hangzhou), China (Shanghai), China (Qingdao), China (Beijing), China (Shenzhen), China (Zhangjiakou), China (Hong Kong), Singapore (Singapore), US (Silicon Valley), or US (Virginia) region.
- The available storage space of the destination ApsaraDB RDS for MySQL instance is larger than the total size of the data in the self-managed TiDB database.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

(?) Note For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database on an ApsaraDB RDS for MySQL instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Pricing.

## **Migration** types

Migration type	Description
	DTS migrates the schemas of required objects to the destination database. DTS supports schema migration for views, tables, and databases.
Schema migration	A warning TiDB and MySQL are heterogeneous databases. DTS does not ensure that the schemas of the source and destination databases are consistent after schema migration. We recommend that you evaluate the impact of data type conversion on your business. For more information, see Data type mappings between heterogeneous databases.
	DTS migrates historical data of required objects to the destination database.
Full data migration	<b>Note</b> During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of the used tablespace of the destination database is larger than that of the source database.

## Preparations

If a whitelist is configured for the TiDB database, you must add the CIDR blocks of DTS servers to the whitelist of the database. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.

1.Configure Sour	rce and Destination	n 💙 2.Configu	ure Migration Types and Objects $ $		3.Advanced Settings	>	4.Precheck
	* Task Name: Ti	iDB_To_RDS					
Course Database							
Source Database				1			
	* Instance Type:	nstance Type: User-Created Database in ECS Instance 🔻			DTS support type		
	* Instance Region:	China (Hangzhou)	Ŧ				
	ECS Instance ID:	i-bp1	•				
	Database Type:     Port Number:	TIDB	•				
*	Database Account:	duttest		1			
*1	Database Password:		<i>ه</i>	] ]	Tect Connectivity 985	sed	
			, , , , , , , , , , , , , , , , , , ,		rest connectivity -		
Destination Database							
	* Instance Type:	RDS Instance	¥				
	* Instance Region:	China (Hangzhou)	T				
	* RDS Instance ID:	rm-bp	-				
*	Database Account:	dtstest					
* [	Database Password:	******	₫>	]	Test Connectivity 🔗 Pas	sed	
	* Encryption:	Non-encrypted OS	SL-encrypted				
						Cano	el Set Whitelist and Next
Section	Paran	neter	Description				
N/A	Task	Name	The task name that DT that you specify a desc task. You do not need	'S a crip to	automatically ge otive name that i specify a unique	enerates. We re makes it easy t e task name.	commend o identify the
			The access method of Created Database in	the 1 EC	e source databa CS Instance is s	se. In this exan selected.	nple, <b>User</b> -
	Instar	nce Type	<b>Note</b> If you set the network environme information, set the set of the set o	ele me ee I	ct other instance nt for the self-n Preparation over	e types, you m nanaged datab rview.	ust deploy base. For
	lnst ar Regio	nce in	The region of the Elast source TiDB database i	tic ( is d	Compute Service leployed.	e (ECS) instance	on which the
	Datab	oase Type	Select T iDB.				
	Port I	Number	The service port numbe 4000.	er o	of the source Til	DB database. D	efault value:
Source Database	Dat ab Accou	oase unt	The account of the sou SELECT permission on t permission.	urce the	e TiDB database objects to migr	. The account i ate and the SH	must have the IOW VIEW

Parameter	Description		
Dat abase Password	The password of the database account. Notice After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		
Instance Type	Select RDS Instance.		
Instance Region	The region where the destination instance resides.		
RDS Instance ID	The ID of the destination instance.		
Database Account	The database account of the destination instance. The account must have read and write permissions on the destination database. For more information about how to create and authorize a database account, see Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account for an ApsaraDB RDS for MySQL instance.		
	The password of the database account.		
Dat abase Password	<b>Notice</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		
Encryption	Select Non-encrypted or SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the ApsaraDB RDS for MySQL instance before you configure the data synchronization task. For more information, see Configure SSL encryption on an ApsaraDB RDS for MySQL instance.		
	ParameterDatabase PasswordInstance TypeInstance RegionRDS Instance IDDatabase AccountDatabase PasswordPatabase Password		

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### □ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration type and the objects to migrate.

Migration	n Types: 🔽 Scher	ma Migration	Full Data Mi	gration			
During ful	ll data migration, d	lata updates in	the source databa	se are not n	nigrated to the destinat	ion instance.	
For data o	consistency, we red	commend that	you select Schema	Migration, I	Full Data Migration, and	Incremental D	Data Migration.
Available					Selected		
Expand the t	tree before you perform	n a glol 🛛 🔍 🔍	<b>^</b>				Q
+ 🚈 dm 🛛	meta				mvsaltest (20bi	ects)	
🗆 📑 myso	qltest				customer		
🖂 📑 Ta	ables				i vipinfo		
🕀 🚈 Vi	iews			>			
🕀 🚈 Fu	unctions			<			
🕀 🚈 Pi	rocedures						
	FORMANCE SCH	FMΔ					
			Ţ				
Select All			Ŧ		Remove All		
Select All Rename Datab	vases and Tables:	Do Not C	•	able Names	Remove All O Change Database and	Table Names	
Select All Rename Datab Retry Time for	vases and Tables: r Failed Connection	Do Not C 720	Change Database and Ta	able Names	Remove All Change Database and	Table Names	
Select All Rename Datab Retry Time for	ases and Tables: r Failed Connection	Do Not C     720	change Database and Ta Minutes (?)	able Names	Remove All Change Database and	Table Names	
Select All Rename Datab Retry Time for <b>nformation:</b> . Data migratio	vases and Tables: r Failed Connection an only copies the data	Do Not C 720 and schema in the	thange Database and Ta Minutes (2)	able Names	Remove All Change Database and Change Database. The destination database. The destination database.	Table Names	not affect any data or sch
Select All Rename Datab Retry Time for <b>nformation:</b> Data migratio the source dal Do not do dal	vases and Tables: r Failed Connection in only copies the data itabase. )L operation during stru	Do Not C     720 and schema in the ucture and full migr	Change Database and Ti Minutes (?) In source database and source d	able Names aves the copy i	Remove All Change Database and in the destination database. T	Table Names	iot affect any data or sch
Select All Rename Datab Retry Time for formation: Data migratio the source dat Do not do DD	ases and Tables: r Failed Connection in only copies the data tabase. NL operation during stru	Do Not C 720 and schema in the ucture and full migr	Change Database and Ta Minutes (7) e source database and s ration, otherwise the tas	able Names aves the copy sk may fail	Remove All Change Database and in the destination database. T	Table Names The process does n	iot affect any data or sch
Select All Rename Datab Retry Time for formation: Data migratio the source da Do not do DD	vases and Tables: r Failed Connection in only copies the data tabase. DL operation during stru	Do Not C     720 and schema in the ucture and full migr	Change Database and Ta Minutes (2) e source database and s ration, otherwise the tag	able Names aves the copy sk may fail	Remove All Change Database and in the destination database. T	Table Names The process does n	not affect any data or sch Save Precheck
Select All Rename Datab Retry Time for <b>nformation:</b> Data migratio the source da Do not do DD	vases and Tables: r Failed Connection in only copies the data (tabase. DL operation during stru Description	Do Not C     720 and schema in the acture and full migr	thange Database and Ta Minutes (?) e source database and s ration, otherwise the tas	able Names aves the copy i ak may fail	Remove All  Change Database and  in the destination database. T  Cancel	Table Names	Not affect any data or sch Save Precheck
Select All Rename Datab Retry Time for <b>tormation:</b> Data migratio the source dal Do not do DD	Pases and Tables: r Failed Connection in only copies the data tabase. DL operation during stru Description	Do Not C     720 and schema in the ucture and full migr	Change Database and Ta Minutes (2) In source database and s ration, otherwise the tas	able Names aves the copy ik may fail	Remove All Change Database and Cancel	Table Names The process does n Previous	tot affect any data or sch Save Precheck
Select All Rename Datab Retry Time for formation: Data migratio the source da Do not do DD	Description	Do Not C     720 and schema in the acture and full migr h Schema	Change Database and Ta Minutes (2) In source database and s ration, otherwise the tas Migration and	able Names aves the copy i sk may fail	Remove All Change Database and The destination database. The Cancel Cancel The Amigration. For	Table Names The process does n Previous Or more info	ot affect any data or sch Save Precheck

the	
migratio	<b>Notice</b> To ensure data consistency, we recommend that you do not write data
n types	to the source TiDB database during data migration.

Setting	Description
Select the objects that you want to migrate	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to add the objects to the Selected section.</li> <li>Notice <ul> <li>You can select columns, tables, or databases as the objects to migrate.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the ApsaraDB RDS instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti ons to the source or destinati on	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
databas e	

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.

- 9. After the task passes the precheck, click **Next**.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

• Warning We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

Task ID/Name:	Status: Completed	View Details   Duplicate Task   Configure Monitoring and Alerting
2019-05-24 17:01:02 Created		2019-05-27 15:54:33 Completed
Schema Migration 100%		Full Data Migration 100%(Migrated Rows: 1000000)
Start Pause Stop Delete		Total: 1 item(s), Per Page: 20 item(s) < < 1 > >

## What's next

After full data migration is complete, you can perform incremental data migration. For more information, see Migrate incremental data from a self-managed TiDB database to an ApsaraDB RDS for MySQL instance.

## 6.8. Source database: Db2

## 6.8.1. Migrate data from a self-managed Db2

## database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a self-managed Db2 database to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a self-managed Db2 database, you can select all of the supported migration types to ensure service continuity.

## Prerequisites

• The version of the Db2 database is 9.7 to 11.5.

**Note** DTS supports data migration from a Db2 for i database of version 7.3 or 7.4 to an ApsaraDB RDS for MySQL instance. You can follow the procedure described in this topic to migrate data from a Db2 for i database to an ApsaraDB RDS for MySQL instance.

• The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total size of the data in the Db2 database.

## Precautions

- In this scenario, DTS cannot synchronize data definition language (DDL) operations.
- If the name of the source database is invalid, you must create a database in the ApsaraDB RDS for MySQL instance before you configure a data migration task.

**Note** For more information about how to create a database and the database naming conventions, see **Create a database on an ApsaraDB RDS for MySQL instance**.

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.
- DTS synchronizes incremental updates from a DB2 database to the destination database based on the Change Data Capture (CDC) replication technology of Db2. However, the CDC replication technology has its own limits. For more information, see General data restrictions for SQL Replication.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## **Migration types**

• Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, index, and foreign key.

• Full data migration

DTS migrates historical data of the required objects from the Db2 database to the destination database in the ApsaraDB RDS for MySQL instance.

• Increment al dat a migration

After full data migration is complete, DTS synchronizes incremental data from the Db2 database to the destination database in the ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data from a Db2 database.

## Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Db2 database	The CONNECT and SELECT permissions	The CONNECT and SELECT permissions	The DBADM permission
ApsaraDB RDS for MySQL instance	The read and write permissions	The read and write permissions	The read and write permissions

For more information about how to create and authorize a database account, see the following topics:

- Db2 database: Creating group and user IDs for a Db2 database installation (Linux and UNIX) and Authorities overview
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

## Data migration process

To prevent data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the Db2 database in the following order:

- 1. Migrate the schemas and indexes.
- 2. Perform full data migration.
- 3. Migrate the schemas of foreign keys.
- 4. Perform incremental data migration.

#### Before you begin

Before you configure an incremental data migration task, enable the archive log feature for the Db2 database. For more information, see Primary log archive method and Secondary log archive method.

**?** Note Skip this step if you perform only full data migration.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

Cancel Set Whitelist and Next

1.Configure Source and Destination Da	atabases 2.Configure Migration Types and Objects	>	3.Map name mo	dification	$\rangle$	4.Precheck
* Task Name:	DB2_TO_RDS					
Source Database						
* Instance Type:	User-Created Database with Public IP Address	٣				
* Instance Region:	Singapore	٣	Get IP Address Segment o	f DTS		
* Database Type:	DB2	۳				
* Hostname or IP Address:	1.00.00					
* Port Number:	50000					
* Database Name:	dtstestdata					
* Database Account:	dtstest					
* Database Password:	•••••	<⊅	Test Connectivity	⊘ Passed		
Destination Database						
* Instance Type:	RDS Instance	٣				
* Instance Region:	Singapore	٣				
* RDS Instance ID:		~				
* Database Account:	dtstest					
* Database Password:	•••••	<⊅	Test Connectivity	⊘ Passed		

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
		Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database with Public IP Address</b> .
	Instance Type	<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .
		If the instance type is set to <b>User-Created Database with Public</b> IP Address, you do not need to specify the instance region.
	Instance Region	<b>?</b> Note If a whitelist is configured for the Db2 database, you must add the CIDR blocks of DTS servers to the whitelist of the database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
	Database Type	Select DB2.
	Hostname or IP Address	Enter the endpoint that is used to connect to the Db2 database. In this example, enter the public IP address.

Section Database	Parameter	Description
		Enter the service port number of the Db2 database. The default port number is <b>50000</b> .
	Port Number	<b>Note</b> The service port of the Db2 database must be accessible over the Internet.
	Database Name	Enter the name of the Db2 database.
	Dat abase Account	Enter the account of the Db2 database. For more information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the Db2 database account.
	Dat abase Password	<b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.
	Instance Type	Select RDS Instance.
	Instance Type Instance Region	Select <b>RDS Instance</b> . Select the region where the destination RDS instance resides.
	Instance Type Instance Region RDS Instance ID	Select RDS Instance.         Select the region where the destination RDS instance resides.         Select the ID of the destination RDS instance.
	Instance Type Instance Region RDS Instance ID Database Account	Select RDS Instance.Select the region where the destination RDS instance resides.Select the ID of the destination RDS instance.Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.
	Instance Type Instance Region RDS Instance ID Database Account	Select RDS Instance.Select the region where the destination RDS instance resides.Select the ID of the destination RDS instance.Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.Enter the password of the database account.
Destinatio n Database	Instance Type Instance Region RDS Instance ID Database Account Database Password	Select RDS Instance. Select the region where the destination RDS instance resides. Select the ID of the destination RDS instance. Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts. Enter the password of the database account. Inter the password of the database account. Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.
Destinatio n Database	Instance Type Instance Region RDS Instance ID Database Account Database Password	Select RDS Instance. Select the region where the destination RDS instance resides. Select the ID of the destination RDS instance. Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts. Enter the password of the database account. Enter the password of the database account. ô Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.
Destinatio n Database	Instance Type Instance Region RDS Instance ID Database Account Database Password	Select RDS Instance. Select the region where the destination RDS instance resides. Select the ID of the destination RDS instance. Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts. Enter the password of the database accounts. Enter the password of the database account. ⑦ Note       After you specify the destination database parameters, click Test Connectivity next to Database         Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.

Section	Parameter	Description
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption on an ApsaraDB RDS for MySQL instance.
		<b>Note</b> The Encryption parameter is available only for regions in mainland China and the China (Hong Kong) region.

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### Q Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

#### Dat a Transmission Service

1.Configu	re Source and    2.Configure Migration Types and   3.Advanced Settings   4.Precheck		
* Migration Ty triggers. For mo	rpes: 🖉 Schema Migration 🖉 Full Data Migration 📝 Incremental Data Migration Note: Incremental data migration does not support re information, see Reference		
Note: do not cleans up the	clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database log too early, the DTS incremental task may fail		
Data migrabi between Aps For long-terr	on applies to snort-term migration scenanos. Typical scenanos include migrating data to the doud, scaling and sharding databases, and migrating data ara Stack databases. n data synchronization in real time, use the data synchronization feature.		
Available	Selected (To edit an object name or its filter, hover over the object and click Edit.) Learn more.		
Expand the	tree before you perform a glol Q		
	ables industrial (20bjects)		
🕀 🚈 V	ews customer		
Select All	Remove All		
* Retry Time fo * Retry Time fo Information: 1. Data migratic in the source dr 2. Do not do Di	passes and Tables: <ul> <li>Do Not Change Database and Table Names</li> <li>Change Database and Table Names</li> <li>r Failed Connection</li> <li>Minutes</li> <li>Minute</li></ul>		
	Cancel Previous Save Precheck		
Setting	Description		
<ul> <li>To perform only full data migration, select Schema Migration and Full Da Migration.</li> </ul>			
Select the	<ul> <li>To ensure service continuity during data migration, select Schema Migration, Full</li> <li>Data Migration, and Incremental Data Migration.</li> </ul>		
migratio n types	<b>Note</b> If <b>Incremental Data Migration</b> is not selected, we recommend that you do not write data to the Db2 database during data migration. This ensures data consistency between the source and destination databases.		

# Data Migration Migrate data from a self-managed database to Alibaba Cloud

Setting	Description			
	Select one or more objects from the <b>Available</b> section and click the <b>&gt;</b> icon to move the objects to the <b>Selected</b> section.			
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination RDS instance, the name of the object remains the same as that in the Db2 database. You can use the object name mapping feature to change the names of the objects that are migrated to the destination RDS instance. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>			
Specify whether to rename object names	You can use the object name mapping feature to change the names of the objects that are migrated to the destination instance. For more information, see Object name mapping.			
Specify the retry time for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.			
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.			

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

1)

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

9. After the data migration task passes the precheck, click Next.

- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the ApsaraDB RDS for MySQL instance.

# 7.Migrate data from a thirdparty cloud to Alibaba Cloud 7.1. Migrate data from an Amazon RDS for MySQL instance to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from an Amazon RDS for MySQL instance to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you configure a data migration task, you can select all of the supported migration types to ensure service continuity.

## Prerequisites

- The **Public accessibility** option of the Amazon RDS for MySQL instance is set to **Yes**. The setting ensures that DTS can access the Amazon RDS for MySQL instance over the Internet.
- An ApsaraDB RDS for MySQL instance is created. For more information, see Create an ApsaraDB RDS for MySQL instance.
- The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total size of the data in the Amazon RDS for MySQL instance.

### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

**Note** For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see **Create a database on an ApsaraDB RDS for MySQL instance**.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.	

## Migration types

#### • Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function. DTS does not support schema migration for events.

#### ? Note

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- DTS does not migrate user information. To call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

#### • Full data migration

DTS migrates historical data of the required objects from the Amazon RDS for MySQL instance to the ApsaraDB RDS for MySQL instance.

#### ? Note

- During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination instance. After full data migration is complete, the tablespace of the destination instance is larger than that of the source instance.
- During schema migration and full data migration, do not perform data definition language (DDL) operations in the source database, for example, add a field. Otherwise, data migration may fail.
- Incremental data migration

After full data migration is complete, DTS retrieves binary log files from the Amazon RDS for MySQL instance. Then, DTS synchronizes incremental data from the Amazon RDS for MySQL instance to the ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data between MySQL databases.

## Permissions required for database accounts
Database	Schema migration	Full data migration	Incremental data migration
Amazon RDS for MySQL	The SELECT permission	The SELECT permission	The REPLICATION CLIENT, REPLICATION SLAVE, SHOW VIEW, and SELECT permissions
ApsaraDB RDS for MySQL	The read and write permissions	The read and write permissions	The read and write permissions

For more information about how to create and authorize a database account, see the following topics:

- Amazon RDS for MySQL instance: Create an account for a user-created MySQL database and configure binary logging
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

## Before you begin

- 1. Log on to the Amazon RDS Management Console.
- 2. Go to the **Basic Information** page of the Amazon RDS for MySQL instance.
- 3. In the **Security group rules** section, click the name of the security group corresponding to the existing inbound rule.

Security group rules (2	2)		C
<b>Q</b> Filter security group rul	es	< 1	0
Security group	▲ Туре	▼ Rule	
	CIDR/IP - Inbound		
100 Con 100 Con 100 Con	CIDR/IP - Outbound	0.0.0/0	

4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

E	EC2 Dashboard	Create Security Group Actions *	∆ ÷ • 0
-	Tags	Q Filter by tags and attributes or search by keyword	② K ≤ 1 to 1 of 1 > >
F	Reports	Name      Group ID     Group Name	VPC ID VPC ID Owner Desc
		default	defau
1	nstances		
1	Edit inbound	rules	×
1	De Type i 3	Protocol (i) Port Range (i) Source (i)	Description (i)
( F	Custom TCP F •	TCP 3306 Custom V	dts 🛛
- 1	Add Rule	4 5 6	
, E	AN NOTE: Any edits m Bu on that rule to be d	de on existing rules will result in the edited rule being deleted and a new rule created opped for a very brief period of time until the new rule can be created.	with the new details. This will cause traffic that depends
			Cancel
3 	Enapshots Lifecycle Manager	Security Group:	
	NETWORK &	Description Inbound Outbound Tags	
: E	Security Groups	Edit 2	
F	Placement Groups	Type (i) Protocol (i) Port Range (i)	Source (i) Description (i)
ł	Key Pairs	All traffic All All	(default)

### ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region. You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- You can add all of the required CIDR blocks to the inbound rule at a time.
- 5. Log on to the Amazon RDS for MySQL database and specify the number of hours to retain binary log files. Skip this step if you do not need to perform incremental data migration.

call mysql.rds\_set\_configuration('binlog retention hours', 24);

#### ? Note

- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of the Amazon RDS for MySQL instance must be enabled and the value of the binlog\_format parameter must be set to row. If the MySQL version is 5.6 or later, the value of the binlog\_row\_image parameter must be set to full.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.

5. Configure the source and destination databases.

1.Configure Source ar	nd Destinatio	n 2.Con	figure Migration Types and Objects	>	3.Map name modific	ration	4.Precheck
* T	ask Name:	IVSOL TO MVSOL					
	Ľ						
Source Database							
* Ir	nstance Type:	User-Created Databa	ase with Public IP Address	•			
* Ins	tance Region:	Singapore		•	Get IP Address Segment o	f DTS	
* Da	atabase Type:	MySQL		*			
* Hostname o	or IP Address:						
*	Port Number:	3306					
* Datal	base Account:	dtstest					
* Databa	ase Password:	•••••		<b>₫</b> >	Test Connectivity	⊘ Passed	
Destination Database							
*1							
• Ir	nstance Type:	RDS Instance		*			
* Ins	tance Region:	Singapore		*			
* RDS	S Instance ID:			•			
* Datab	Dase Account:	dtstest		4	Test Conservitivity		
Databa	ise Password.	•••••		ą,	Test Connectivity	C Passeu	
							Cancel Set Whitelist and Next
Section	Parar	meter	Description				
N/A	Task	Name	DTS automatical specify an inform to use a unique t	ly gene native r ask na	erates a task n name for easy me.	ame. We rec identificatio	commend that you n. You do not need
	Insta	nce Type	Select User-Crea	ated D	)atabase wit	h Public IP	Address.
	lnsta Regio	nce on	lf the instance ty IP Address, you	pe is so do no	et to <b>User-Cr</b> o t need to spec	eated Data cify the insta	base with Public ance region.
	Data	base Type	Select MySQL.				

Section	Parameter	Description					
Section	Hostname or IP Address	Enter the endpoint that is used to access the Amazon RDS for MySQL instance.					
	Port Number	Enter the service port number of the Amazon RDS for MySQL instance. The default port number is <b>3306</b> .					
	Dat abase Account	Enter the database account of the Amazon RDS for MySQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.					
	Dat abase Password	Enter the password of the database account. <b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.					
	Instance Type	Select RDS Instance.					
	Instance Region	Select the region where the ApsaraDB RDS for MySQL instance resides.					
	RDS Instance ID	Select the ID of the ApsaraDB RDS for MySQL instance.					
	Dat abase Account	Enter the database account of the ApsaraDB RDS for MySQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.					

Section	Parameter	Description
Destinatio n Database		Enter the password of the database account.
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption on an ApsaraDB RDS for MySQL instance
		<b>Note</b> The <b>Encryption</b> parameter is available only for regions in mainland China and the China (Hong Kong) region.

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗋 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

### Dat a Transmission Service

# Data Migration Migrate data from a third-party cloud to Alibaba Cloud

1.Configure Source and     Migration Types: Schema Migration triggers. For more information, see Referen Note: do not clean up the incremental di cleans up the log too early, the DTS incr Data migration applies to short-term mig between Apsara Stack databases. For long-term data synchronization in re Available	Z.Configure Migration Type     ✓ Full Data Migration     ree     ata log generated by the source of     emental task may fail     gration scenarios. Typical scenario     al time, use the data synchroniza	Incremental D The Incremental D database after the as include migratin tion feature.	3. Advanced St Data Migration P DTS task is started of g data to the cloud, Selected (To edi	Note: Incremental data mi when the DTS full task is r scaling and sharding data	4.Precneck gration does not support unning. If the source database bases, and migrating data er, hover over the object and click
Expand the tree before you perform a g	lo Q	> <	Edt.) Learn mor	e. ata (20bjects) er	
Select All			Remove All		
*Rename Databases and Tables: * Retry Time for Failed Connection *Source table DMS_ ONLINE_ Do you want to copy the temporary table to the target database during DDL: Information: 1. Data migration only copies the data and in the source database. 2. Do not do DDL operation during structur	Do Not Change Database a     720     Minutes     Yes     No     O     res     and full migration, otherwise th	nd Table Names	<ul> <li>Change Data</li> <li>in the destination d</li> </ul>	base and Table Names atabase. The process does	s not affect any data or schema
				Cancel Previous	Save Precheck

# Data Migration Migrate data from a third-party cloud to Alibaba Cloud

Setting	Description
	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> </ul>
Select the migratio n types	<ul> <li>Note</li> <li>If Incremental Data Migration is not selected, we recommend that you do not write data to the source database during data migration. This ensures data consistency between the source and destination databases.</li> <li>During schema migration and full data migration, we recommend that you do not perform DDL operations on the required objects. Otherwise, the objects may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

Setting	Description
Specify whether to copy tempora ry tables to the destinati on databas e when DMS perform s online DDL operatio ns on the source table	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.
	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
	<b>Note</b> If you select No, the tables in the destination database may be locked.

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the ApsaraDB RDS for MySQL instance.

# 7.2. Migrate data from an Amazon RDS for Oracle instance to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from an Amazon RDS for Oracle instance to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you configure a data migration task, you can select all of the supported migration types to ensure service continuity.

## Prerequisites

- The **Public accessibility** option of the Amazon RDS for Oracle instance is set to **Yes**. The setting ensures that DTS can access the Amazon RDS for Oracle instance over the Internet.
- The database version of the Amazon RDS for Oracle instance is 9i, 10g, 11g, or 12c or later (nonmultitenant architecture).
- The database version of the ApsaraDB RDS for MySQL instance is 5.6 or 5.7.
- The available storage space of the ApsaraDB RDS for MySQL instance is at least twice the total size of the data in the Amazon RDS for Oracle instance.

Onte The binary log files that are generated during data migration occupy some space. They are automatically cleared after data migration is complete.

## Precautions

• DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no

primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.

- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

**?** Note For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database on an ApsaraDB RDS for MySQL instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# Migration types

• Schema migration

DTS supports schema migration for the following types of objects: table, index, constraint, and sequence. DTS does not support schema migration for the following types of objects: view, synonym, trigger, stored procedure, function, package, and user-defined type.

• Full dat a migration

DTS migrates historical data of the required objects from the source database in the Amazon RDS for Oracle instance to the destination database in the ApsaraDB RDS for MySQL instance.

• Incremental data migration

DTS retrieves redo log files from the source database in the Amazon RDS for Oracle instance. Then, DTS synchronizes incremental data from the source database in the Amazon RDS for Oracle instance to the destination database in the ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data from an Oracle database.

# SQL operations that can be synchronized during incremental data migration

• INSERT, DELETE, and UPDATE

#### • CREATE TABLE

**?** Note If a CREATE TABLE operation creates a partitioned table or a table that contains functions, DTS does not synchronize the operation.

- ALTER TABLE, including only ADD COLUMN, DROP COLUMN, RENAME COLUMN, and ADD INDEX
- DROP TABLE
- RENAME TABLE, TRUNCATE TABLE, and CREATE INDEX

## Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Amazon RDS for Oracle instance	The permissions of the schema owner	The permissions of the schema owner	The permissions of the master user
ApsaraDB RDS for MySQL instance	The read and write permissions on the destination database	The read and write permissions on the destination database	The read and write permissions on the destination database

For more information about how to create and authorize database account, see the following topics:

- Amazon RDS for Oracle instance: CREATE USER and GRANT
- ApsaraDB RDS for MySQL instance: Create accounts and databases for an ApsaraDB RDS for MySQL instance

## Data type mappings

For more information, see Data type mappings between heterogeneous databases.

## Before you begin

- 1. Log on to the Amazon RDS Management Console.
- 2. Go to the Basic Information page of the Amazon RDS for Oracle instance.
- 3. In the **Security group rules** section, click the name of the security group corresponding to the existing inbound rule.

Security group rules	(2)		C
<b>Q</b> Filter security group	rules	< 1	> ©
Security group	▲ Туре	▼ Rule	•
	CIDR/IP - Inbound		
101-101-101-101-101-101-101-101-101-101	CIDR/IP - Outbound	0.0.0/0	

4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises

#### databases.

	EC2 Da Events	ashboard		Create Security (	Group Actions 🛩						4 <del>0</del>	ø	Ø
	Tags			Q, Filter by tags a	ind attributes or search b	y keyword				0 K < 1	to 1 of	1 >	
	Report	S		Name	<ul> <li>Group ID</li> </ul>	▲ Gro	oup Name	· VPC ID		Owner		•	Desc
	Limits				1000	defa	ult	10000					defau
	Instanc	ces es											
l	La										_		
	Sp	Edit inbou	ind ru	lles							2	×	
l	De	Type () 3		Protocol (j)	Port Range (j)	Source ()			Description (i)				
l	Ca	Custom TCP	FT	TCP	1521	Custom •			dts		⊗		
	IM	Add Rule			4	5	6						
l	AN	NOTE: Any edit	s made	on existing rules wil	Il result in the edited ru	le being deleted a	nd a new rule create	ed with the new deta	ils. This will cause tra	affic that depend	is		
l	Bu	on that rule to b	be dropp	bed for a very brief	period of time until the	new rule can be c	reated.				7		
-	ST									Cancel	Save	1	
l	Vo	nots											•
l	Lifecyc	le Manager		Security Group:							_		Î
	NETWO	IRK &		Description	nbound Outboun	d Tags							
ī	Securi	ty Groups		Edit 2									
l	Elastic	IPs											
	Placem	ent Groups		Туре (і)	Protoco		Port Range (i)	Sourc	e (i)	Description	()		
	кеу Ра	irs	¥	All traffic	All		All	100	(default)				Ŧ

#### ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region. You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- You can add all of the required CIDR blocks to the inbound rule at a time.
- 5. Modify the log settings of the Amazon RDS for Oracle instance. Skip this step if you do not need to perform incremental data migration.
  - If the database version of the Amazon RDS for Oracle instance is 12c or later (non-multitenant architecture), perform the following steps to configure the log settings:
    - a. Use the master user account and the SQL\*Plus tool to connect to the Amazon RDS for Oracle instance.
    - b. Enable archive logging and supplement al logging.

Туре

Procedure

Туре	Procedure		
	a. Execute the following statement to check whether archive logging is enabled:		
	SELECT LOG_MODE FROM v\$database;		
	b. Set a retention period for archived logs.		
Archive logging	<ul> <li>Note We recommend that you set the retention period of archived logs to at least 72 hours.</li> <li>In this example, the retention period is set to 72 hours.</li> </ul>		
	<pre>exec rdsadmin.rdsadmin_util.show_configuration; exec rdsadmin.rdsadmin_util.set_configuration('arc hivelog retention hours', 72);</pre>		

Туре	Procedure				
	<ul> <li>Enable supplemental logging at the database or table level:</li> <li>Enable database-level supplemental logging</li> <li>a. Execute the following statement to check whether database-level supplemental logging is enabled:</li> </ul>				
	SELECT supplemental_log_data_min, supplemental_log_data_pk, supplemental_log_data_ui FROM v\$database;				
	b. Enable primary key and unique key supplemental logging at the database level:				
	<pre>exec rdsadmin.rdsadmin_util.alter_supplemental_l ogging('ADD', 'PRIMARY KEY'); exec rdsadmin.rdsadmin_util.alter_supplemental_l ogging('ADD', 'UNIQUE');</pre>				
Supplemental logging	<ul> <li>Enable table-level supplemental logging by using one of the following methods:</li> <li>Enable table-level supplemental logging for all columns:</li> </ul>				
	<pre>exec rdsadmin.rdsadmin_util.alter_supplemental_log ging('ADD', 'ALL');</pre>				
	<ul> <li>Enable primary key supplemental logging at the table level:</li> </ul>				
	<pre>exec rdsadmin.rdsadmin_util.alter_supplemental_log ging('ADD', 'PRIMARY KEY');</pre>				

- c. Grant fine-grained permissions to the database account of the Amazon RDS for Oracle instance.
  - Grant fine-grained permissions
- If the database version of the Amazon RDS for Oracle instance is 9i, 10g, or 11g, perform the following steps to configure the log settings:
  - a. Use the master user account and the SQL\*Plus tool to connect to the Amazon RDS for

Oracle instance.

b. Run the archive log list; command to check whether the Amazon RDS for Oracle instance is running in ARCHIVELOG mode.

**Note** If the instance is running in NOARCHIVELOG mode, switch the mode to ARCHIVELOG. For more information, see Managing Archived Redo Logs.

#### c. Enable force logging.

exec rdsadmin.rdsadmin\_util.force\_logging(p\_enable => true);

#### d. Enable supplement al logging for primary keys.

```
begin rdsadmin.rdsadmin_util.alter_supplemental_logging(p_action => 'ADD',p_type
=> 'PRIMARY KEY');end;/
```

#### e. Enable supplement al logging for unique keys.

```
begin rdsadmin.rdsadmin_util.alter_supplemental_logging(p_action => 'ADD',p_type
=> 'UNIQUE');end;/
```

#### f. Set a retention period for archived logs.

```
begin rdsadmin.rdsadmin_util.set_configuration(name => 'archivelog retention hour
s', value => '24');end;/
```

#### g. Set a retention period for archived logs.

**Note** We recommend that you set the retention period of archived logs to at least 24 hours.

#### h. Commit the changes.

commit;

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.

1 Configure Source and Des	tination Databases 2 Confi	nure Migration Types and Objects	3 Man name modification	4 Prochack
				In recieck
* Task	Name: Oracle_TO_RDS			
Source Database				
* Insta	nce Type: User-Created Database w	th Public IP Address 🔹 🕻	DTS support type	
* Instanc	e Region: China (Hangzhou)	<b>T</b>	Get IP Address Segment of DTS	
* Datab	ase Type: Oracle	<b>T</b>		
* Hostname or IF	Address:			
* Port	t Number: 1521			
Insta	* CID:	AC Instance		
* Database	Account:			
* Database	Password:		Test Connectivity	
		*	rest connectivity	
Destination Database				
* Insta	nce rype: RDS Instance	•		
* Instanc	Le Region: China (Hangzhou)	•		
* RDS Ins	stance ID:	•		
* Database	Account:			
* Database i	Password:	♥	Test Connectivity	
		and ypecu		
			Cancel Assess Data Migration	n to Cloud Set Whitelist and Next
Section	Parameter	Description		
N/A	Task Name	DTS automatically g specify an informati to use a unique task	jenerates a task name. We recon we name for easy identification. c name.	וmend that you You do not need
	Instance Type	Select User-Create	d Database with Public IP Ac	ldress.
	Instance Region	If you select <b>User-C</b> the instance type, yo parameter.	<b>reated Database with Public</b> ou do not need to specify the <b>In</b>	IP Address as stance Region
	Database Type	Select Oracle.		

Section	Parameter	Description					
		Enter the endpoint that is used to access the Amazon RDS for Oracle instance.           O         Note         You can obtain the endpoint on the Basic           Information page of the Amazon RDS for Oracle instance.					
Source Database	Hostname or IP Address	Amazon RDS       X       Connectivity & security       Monitoring       Lags & events       Configuration       Maintenance & backups         Dashboard       Tags       Tags					
	Port Number	Enter the service port number of the Amazon RDS for Oracle instance. The default port number is <b>1521</b> .					
	<ul> <li>If you select Non-RAC Instance, you must specify the SID parameter.</li> <li>Instance Type</li> <li>If you select RAC Instance, you must specify the Service Name parameter.</li> <li>In this example, select Non-RAC Instance and specify the SID.</li> </ul>						
	Dat abase Account	Enter the database account of the Amazon RDS for Oracle instance. For information about the permissions that are required for the account, see Permissions required for database accounts.					
	Dat abase Password	Enter the password of the database account. <b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.					
	Instance Type	Select RDS Instance.					
	Instance Region	Select the region where the ApsaraDB RDS for MySQL instance resides.					
	RDS Instance ID	Select the ID of the ApsaraDB RDS for MySQL instance.					

Section	Parameter	Description				
Destinatio n	Database Account	Enter the database account of the ApsaraDB RDS for MySQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.				
Database		Enter the password of the database account.				
	Dat abase Password	<b>Note</b> After you specify the destination database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.				

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### Q Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

#### Dat a Transmission Service

### Dat a Migration Migrate dat a from a

third-party cloud to Alibaba Cloud

1.Configure So	purce and Destination > 2.Configure Migration Types and > 3.Advanced Settings > 4.Precheck									
* Migration triggers, Fo	ion Types: 🗹 Schema Migration 🖉 Full Data Migration 📝 Incremental Data Migration Note: Incremental data migration does not support									
Note: de cleans u	Vote: do not clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database cleans up the log too early, the DTS incremental task may fail									
Available	e Selected (To edit an object name or its filter, hover over the object and click									
Expand	d the tree before you perform a gloi   Q									
🗉 🦢 te	estdb									
	>									
	<									
Select All										
*Panama I	Remove All									
* Retry Tir	me for Failed Connection 720 Minutes 2									
<b>Informati</b> 1. Data mi	ion: igration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema									
in the sour 2. Do not o	rce database. do DDL operation during structure and full migration, otherwise the task may fail									
	Cancel Previous Save Prechee									
etting	Description									
	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration</li> </ul>									
elect	<ul> <li>To ensure service continuity during data migration, select Schema Migration, Full</li> <li>Data Migration, and Incremental Data Migration.</li> </ul>									
iigratio types	<b>Note</b> If <b>Incremental Data Migration</b> is not selected, we recommend that you do not write data to the source database during data migration. This ensures data									

Setting	Description				
	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.				
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the ApsaraDB RDS for MySQL instance. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>				
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.				
Specify the retry time for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.				
the source or destinati on databas e	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.				

#### 8. In the lower-right corner of the page, click **Precheck**.

## ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the *icon next* to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.

- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

**Notice** We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the ApsaraDB RDS for MySQL instance.

# 7.3. Migrate incremental data from an Amazon RDS for PostgreSQL instance to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to migrate incremental data from an Amazon RDS for PostgreSQL instance to an ApsaraDB RDS for PostgreSQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. You can select all of the supported migration types to ensure service continuity.

## Prerequisites

- The version of the Amazon RDS for PostgreSQL instance is 10.4 to 12.
- The **Public accessibility** option of the Amazon RDS for PostgreSQL instance is set to **Yes**. This ensures that DTS can access the instance over the Internet.
- The value of the rds.logical\_replication parameter is set to *1*. This ensures that DTS can read increment al data from the Amazon RDS for PostgreSQL instance.

• An ApsaraDB RDS for Post greSQL instance is created. For more information, see Create an ApsaraDB RDS for Post greSQL instance.

### ? Note

- The version of the ApsaraDB RDS for PostgreSQL instance is 10 or 11. To migrate data between different database versions, create a pay-as-you-go instance to verify compatibility.
- The available storage space of the ApsaraDB RDS for PostgreSQL instance must be larger than the total size of the data in the Amazon RDS for PostgreSQL instance.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The objects to be migrated must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records and data migration may fail.
- A single data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each database.
- In this scenario, DTS can migrate only data manipulation language (DML) operations, such as INSERT, DELETE, and UPDATE.
- During data migration, DTS creates a replication slot for the Amazon RDS for PostgreSQL instance. The replication slot is prefixed with dts\_sync\_. DTS automatically clears historical replication slots every 90 minutes to reduce storage usage.

(?) Note If the data migration task is released or fails, DTS automatically clears the replication slot. If a primary/secondary switchover is performed on the Amazon RDS for PostgreSQL instance, you must log on to the secondary database to clear the replication slot.

Qu	Query Editor Query History					Scratch Pad			
1	<pre>1 SELECT * FROM pg_replication_slots;</pre>								
Da	ta Output	Explain Messa	ges Notifica	ations					
	slot_name name	<u></u>	plugin name	slot_type text	datoid oid	database name	temporary boolean	active boolean	active_pid integer
1	dts_sync_ol	าน	pgoutput	logical	16	dtstestdata	false	true	

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## Billing

<sup>&</sup>gt; Document Version: 20220712

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# Migration types

• Schema migration

DTS migrates the schemas of the required objects to the ApsaraDB RDS for PostgreSQL instance. DTS supports schema migration for the following types of objects: table, trigger, view, sequence, function, user-defined type, rule, domain, operation, and aggregate.

⑦ Note DTS does not migrate functions that are written in the C programming language.

• Full data migration

DTS migrates historical data of the required objects from the Amazon RDS for PostgreSQL instance to the ApsaraDB RDS for PostgreSQL instance.

• Incremental data migration

After full data migration is complete, DTS synchronizes incremental data from the Amazon RDS for PostgreSQL instance to the ApsaraDB RDS for PostgreSQL instance. Incremental data migration allows you to ensure service continuity when you migrate data between PostgreSQL databases.

## Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Amazon RDS for PostgreSQL	The USAGE permission on pg_catalog	The SELECT permission on the objects to be migrated	The rds_superuser permission
ApsaraDB RDS for PostgreSQL	The CREATE and USAGE permissions on the objects to be migrated	The permissions of the schema owner	The permissions of the schema owner

## Data migration process

To prevent data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

- 1. Migrate the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates.
- 2. Perform full data migration.
- 3. Migrate the schemas of triggers and foreign keys.
- 4. Perform incremental data migration.

**?** Note Before incremental data migration, do not perform data definition language (DDL) operations on the objects in the Amazon RDS for PostgreSQL instance. Otherwise, the objects may fail to be migrated.

# Before you begin

- 1. Log on to the Amazon RDS Management Console.
- 2. In the upper-right corner of the page, select the region where the destination instance resides.
- 3. In the left-side navigation pane, click **Databases**. On the page that appears, click the ID of the destination database. The **Basic Information** page appears.

Amazon RDS ×	RDS > Databases
Dashboard Databases Performance Insights	Databases C Group resources C Modify Actions Restore from S3 Create database
Snapshots Automated backups	Q Filter databases < 1 > @
Reserved instances	→ DB identifier AZ ▼
Subnet groups Parameter groups	Instance     PostgreSQL     ap-southeast-1c

4. In the **Security group rules** section, click the name of the security group corresponding to the existing inbound rule.

Security group rules (2)			C
<b>Q</b> Filter security group rules		< 1 >	0
Security group	Туре	▼ Rule	
	CIDR/IP - Inbound		
MARK DO MARK	CIDR/IP - Outbound	0.0.0.0/0	

5. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

	EC2 Dashboard	Create Security Group	Actions 👻			<u>ح</u>	e •	0
	Tags	Q, Filter by tags and attri	butes or search by keyword				1 of 1 →	
	Reports	Name - C	Group ID	Group Name	VPC ID	• Owner	-	Desc
•	INSTANCES			default	0.000	-		defau
	Instances							
	sp Edit inbound	rules					×	
	Re De Type () 3	Protocol (j Port	Range (i) Source	1)	Descripti	on (i)		
	Ca Re	TCP 5432	Custom	•	dts		8	
•	IM. Add Rule		4 5	6				
	AN NOTE: Any edits made Bu on that rule to be dro	de on existing rules will result opped for a very brief period	t in the edited rule being de of time until the new rule ca	leted and a new rule created n be created.	with the new details. This will ca	ause traffic that depends		
•	EL					Cancel	ve	
	Vo							Þ
	Lifecycle Manager	Security Group:						
=	NETWORK & SECURITY	Description	d Outbound Tags					
I	Security Groups	Edit 2						
	Placement Groups	Туре ()	Protocol (j)	Port Range (j)	Source (j)	Description ()		
	Key Pairs	All traffic	All	All	(defau	ult)		•

### ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region. You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- You can add all of the required CIDR blocks to the inbound rule at a time.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

### Dat a Transmission Service

1.Set Source and Desti	nation Databases	2.Mig	ration Types and Tasks		3.Precheck	
* Task Name:	10000					
Source Database						
* Instance Type:	User-Created Database with Public IP Address	•				
Source Instance Region:	Singapore	٣	Obtain IP Address Segment of DTS			
* Database Type:	PostgreSQL	٣				
* Hostname or IP Address:						
* Port Number:	5432					
* Database Name:	dtstestdata					
* Database Account:						
* Database Password:	•••••	<⊅	Test Connection			
Destination Database						
Describer of Database						
* Instance Type:	RDS Instance	۳				
* Source Instance Region:	China (Hangzhou)	*				
* RDS Instance ID:		•				
* Database Name:	dtstestdata					
* Database Account:						
* Database Password:	•••••	<⊅	Test Connection			
				Cancel	Assess Migration	Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select User-Created Database with Public IP Address.
	Instance Region	Select the region where the source instance resides. If you select <b>User-Created Database with Public IP Address</b> as the instance type, you do not need to specify the Instance Region parameter.
	Database Type	Select PostgreSQL.

Section	Parameter	Description					
Source Dat abase	Hostname or IP Address	Enter the endpoint that is used to access the Amazon RDS for bostgreSQL instance.					
	Port Number	Enter the service port number of the Amazon RDS for PostgreSQL instance. The default port number is <b>5432</b> .					
	Database Name	Enter the name of the source database in the Amazon RDS for PostgreSQL instance.					
	Dat abase Account	Enter the database account of the Amazon RDS for PostgreSQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.					
	Dat abase Password	Enter the password of the database account. <b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.					
	Instance Type	Select RDS Instance.					

Section	Parameter	Description					
	Instance Region	Select the region where the ApsaraDB RDS for PostgreSQL instance resides.					
	RDS Instance ID	Select the ID of the ApsaraDB RDS for PostgreSQL instance.					
		Enter the name of the destination database in the ApsaraDB RDS for PostgreSQL instance. The name can be different from the name of the source database in the Amazon RDS for PostgreSQL instance.					
Destination Database	Database Name	<b>Note</b> Before you configure the data migration task, you must create a database in the ApsaraDB RDS for PostgreSQL instance. For more information, see <b>Create a database on an ApsaraDB RDS for PostgreSQL instance</b> .					
	Database Account	Enter the database account of the ApsaraDB RDS for PostgreSQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.					
		Enter the password of the database account.					
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.					

### 6. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note DTS adds the CIDR blocks of DTS servers to the whitelist of the ApsaraDB RDS for PostgreSQL instance. This ensures that DTS servers can connect to the ApsaraDB RDS for PostgreSQL instance.

#### 7. Select the migration types and the objects to be migrated.

#### Dat a Transmission Service

### Data Migration• Migrate data from a

third-party cloud to Alibaba Cloud

1.Configure Source and	Destination	2.Configure Migra	tion Types and	3.Advanced Settings	$\rightarrow$	4.Precheck
<ul> <li>Migration Types: triggers. For more in</li> </ul>	Schema Migration	n 🔽 Full Data Mi ance	igration 🔽 Incremental	Data Migration Note: Increme	ental data migrati	on does not support
Note: do not clea cleans up the log	n up the incremental too early, the DTS inc	data log generated by rremental task may fail	the source database after th I	e DTS task is started when the DTS	5 full task is runni	ng. If the source database
Available				Selected (To edit an object na Edit.) Learn more.	ame or its filter, h	over over the object and click
Expand the tree	before you perform a	glol Q				Q
🗈 🥁 test123				📔 dtstestdata		
			>			
			<			
Select All						
Select All				Remove All		
*Rename Databases	s and Tables:	Do Not Change	Database and Table Names	<ul> <li>Change Database and Tab</li> </ul>	de Names	
* Retry Time for Fail	led Connection	720	Minutes (?)			
Information: 1. Data migration on in the source databa 2. Do not do DDL or	nly copies the data and use. peration during structu	d schema in the source	e database and saves the cop otherwise the task may fail	y in the destination database. The	process does not	t affect any data or schema
	-					
				Cancel	Previous	Save Precheck
etting	Descriptio	n				
elect the	Select <b>Sc</b> l	hema Migra	tion Full Data	Migration and In	crement	al Data Migratio

Setting	Description
Select the objects to be migrated	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to add the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or schemas as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination instance, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the ApsaraDB RDS for PostgreSQL instance. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connections to the source or destination database	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>(?)</b> Note When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

## 8. In the lower-right corner of the page, click **Precheck**.

? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click **Next**.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

(?) Note A task does not automatically stop during incremental data migration. You must manually stop the task. We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination instance.

- i. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- ii. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the ApsaraDB RDS for PostgreSQL instance.

# 7.4. Migrate full data from an Amazon RDS for PostgreSQL instance to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to migrate full data from an Amazon RDS for PostgreSQL instance to an ApsaraDB RDS for PostgreSQL instance by using Data Transmission Service (DTS).

## Prerequisites

- An Amazon RDS for Post greSQL instance is created and the database engine version is 9.4, 9.5, 9.6, or 10.0.
- The **Public accessibility** option of the Amazon RDS for PostgreSQL instance is set to **Yes**. This ensures that DTS can access the instance over the Internet.
- An ApsaraDB RDS for PostgreSQL instance is created. For more information, see Create an ApsaraDB RDS for PostgreSQL instance.

Onte The database engine version of the ApsaraDB RDS for PostgreSQL instance is 9.4 or 10.0.

• The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total size of the data in the Amazon RDS for PostgreSQL instance.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- Incremental data migration is not supported in this scenario. Before you start the data migration task, you must stop the services that run on the Amazon RDS for PostgreSQL instance. To ensure data consistency, we recommend that you do not write data to the Amazon RDS for PostgreSQL instance during data migration.
- Each data migration task can migrate data from only a single database. To migrate data from multiple databases, you must create a data migration task for each database.
- Functions that are written in the C programming language cannot be migrated.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If a data migration task fails, DTS automatically resumes the task. Therefore, before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance overwrites the data in the destination instance after the task is resumed.
- To ensure that the data migration task runs as expected, you can perform a primary/secondary switchover only on a V11 ApsaraDB RDS for PostgreSQL instance. In this case, you must set the rds\_failover\_slot\_mode parameter to sync . For more information, see Logical Replication Slot Failover.

**Warning** If you perform a primary/secondary switchover on a self-managed PostgreSQL database or an ApsaraDB RDS for PostgreSQL instance of a version other than V11, the data migration task stops.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Pricing.

# **Migration types**

• Schema migration

DTS migrates the schemas of required objects to the ApsaraDB RDS for PostgreSQL instance. DTS supports the following types of objects for schema migration: table, trigger, view, sequence, function, user-defined type, rule, domain, operation, and aggregate.

• Full dat a migration

DTS migrates the historical data of required objects from the Amazon RDS for PostgreSQL instance to the ApsaraDB RDS for PostgreSQL instance.

## Permissions required for database accounts

Database	Schema migration	Full data migration
Amazon RDS PostgreSQL	The USAGE permission on pg_catalog	The SELECT permission on the objects to migrate
ApsaraDB RDS for PostgreSQL	The CREATE and USAGE permissions on the objects to migrate	Permissions of the schema owner

## Process of full data migration

To prevent data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

- 1. Migrate the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates.
- 2. Perform full data migration.
- 3. Migrate the schemas of triggers and foreign keys.

# Preparation 1: Edit the inbound rule of the Amazon RDS for PostgreSQL instance

- 1. Log on to the Amazon RDS Management Console.
- 2. Go to the **Basic Information** page of the Amazon RDS for PostgreSQL instance.
- 3. In the **Security group rules** section, click the name of the security group corresponding to the existing inbound rule.

Security group rules (2)	C						
<b>Q</b> Filter security group rules	<b>Q</b> Filter security group rules						
Security group	Туре	•	Rule 🔻				
	CIDR/IP - Inbound						
MARKET OF MILLION	CIDR/IP - Outbound		0.0.0/0				

4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

	EC2 D	ashboard		Create Security (	Group Actions 👻						<b>∆</b> ∂	ø	0
	Tags	,		Q, Filter by tags a	nd attributes or search b	y keyword			0		1 to 1 of 1		
	Report	ts		Name	• Group ID	▲ Gi	oup Name	VPC ID	- Ow	ner		•	Desc
-	INSTAN	ICES			1000	de	fault	10.000					defau
	Instand	ces											
	Sp	Edit inbound	d ru	iles							>	<	
	De	Type (j 3	F	Protocol (i)	Port Range (j)	Source (i)			Description (i)				
	Ca Re	Custom TCP F 🔻		TCP	5432	Custom •			dts		⊗		
		Add Rule			4	5	6						
	Bu	NOTE: Any edits m on that rule to be d	iade Iropp	on existing rules wil ed for a very brief	I result in the edited ru period of time until the	new rule can be	and a new rule created created.	d with the new deta	ails. This will cause traffic	that depen	ids		
	EL/ ST									Cancel	Save		
l	Snaps	hots		Security Group:	-								
-	Lifecyo NETWO SECUR	Cle Manager		Description	nbound Outboun	d Tags							
L	Secur Elastic	rity Groups		Edit 2									
	Placen	nent Groups		Туре ()	Protoce	ol (j)	Port Range (j)	Sour	ce (j)	escriptio	n (j		
	Key Pa	airs *		All traffic	All		All	100	(default)				-

## ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region.
   You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- $\circ~$  You can add all of the required CIDR blocks to the inbound rule at a time.

# Preparation 2: Create a database and schema in the destination ApsaraDB RDS instance

Create a database and schema in the destination ApsaraDB RDS instance based on the database and schema information of the objects to migrate. The schema name of the source and destination databases must be the same. For more information, see Create a database on an ApsaraDB RDS for PostgreSQL instance and Appendix: User and schema management.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Set Source and D	Destination Databases	2.Mi	igration Types and Tasks	>	3.Precheck	
* Task Name:						
Source Database						
* Instance Typ	pe: User-Created Database with Publ	ic IP Address 🔻				
* Source Instance Region	Singapore		Obtain IP Address Segment of DTS			
<ul> <li>Database Typ</li> </ul>	PostgreSQL					
* Hostname or IP Addre	per a se la propie de la reseau de					
* Port Numb	5432					
* Database Nan	dtstestdata					
* Database Accou	int:					
Database Password:		4>	Test Connection			
Dectination Database						
Desnuerou Derensee						
* Instance Type: RDS Instance		•				
<ul> <li>Source Instance Regi</li> </ul>	on: China (Hangzhou)	•				
* RDS Instance		•				
* Database Accou	Int: dtstestdata					
* Database Passwo	rd:	đs	Tast Connection			
		Ţ	rest connection			
					Cancel Assess Migration Set Whitelist and Next	
Section	Parameter	Description				
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.				
	Instance Type	Select User	-Created Databa	ise with Pul	blic IP Address.	
	Instance Region	The region where the source instance resides. If you select <b>User-</b> <b>Created Database with Public IP Address</b> for the instance type, you do not need to specify the Instance Region parameter.				
	Database Type	Select Post	greSQL.			

Section	Parameter	Description					
Source Instance Details	Hostname or IP Address	The endpoint that is used to access the Amazon RDS for PostgreSQL instance.					
		<ul> <li>Note You can find the endpoint on the Basic</li> <li>Information page of the Amazon RDS for PostgreSQL instance.</li> </ul>					
		Amazon RDS     X     Connectivity & security     Monitoring     Logs & events     Configuration     Maintenance & backups       Dashboard     Tags     Tags     Tags     Tags     Tags     Tags					
		Performance Insights Snapshots Automated backups					
		Reserved instances     Endpoint & port     Networking     Security       Submet groups     Endpoint     Availability zone     VPC security groups       Parameter groups     Port     VPC       Option groups     Port     VPC       Events     Submet group     VPC       Events     Submet group     VPC       Events     Submet group     VPC       Events     Submet group     Certificate authority       Recommendations ()     Submet group     Certificate authority date       Mar Geb, 2020     Mar Geb, 2020     Mar Geb, 2020					
	Port Number	The service port number of the Amazon RDS for PostgreSQL instance. Default value: <b>5432</b> .					
	Database Name	The name of the source database in the Amazon RDS for PostgreSQL instance.					
	Dat abase Account	The database account of the Amazon RDS for PostgreSQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.					
	Dat abase Password	The password of the database account.					
		<b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.					
	Instance Type	Select RDS Instance.					
	Instance Region	The region where the ApsaraDB RDS for PostgreSQL instance resides.					
	RDS Instance ID	The ID of the ApsaraDB RDS for PostgreSQL instance.					
Parameter	Description						
-----------------------	--	--	--	--			
	The name of the destination database in the ApsaraDB RDS for PostgreSQL instance. The name can be different from the name of the source database in the Amazon RDS for PostgreSQL instance.						
Database Name	Note Before you configure the data migration task, you must create a database and schema in the ApsaraDB RDS for PostgreSQL instance. For more information, see Preparation 2: Create a database and schema in the destination ApsaraDB RDS instance.						
Dat abase Account	The database account of the ApsaraDB RDS for PostgreSQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.						
	The password of the database account.						
Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.						
	Dat abase Name Dat abase Account Dat abase Password						

6. In the lower-right corner of the page, click Set Whitelist and Next.

### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select migration types and the objects to migrate.

<ul> <li>Migration Types: Schema Migration Full Data Migration Incremental Data Migration</li> <li>Data migration applies to short-term migration scenarios. Typical scenarios include migrating data to the cloud, scaling and she between Apsara Stack databases.</li> <li>For long-term data synchronization in real time, use the data synchronization feature.</li> </ul>	arding databases, and migrating data
Available       Selected (To edit an object nam         If you search globally, please expand the       Q         If you search globally, please expand the       Q	e or its filter, hover over the object and dick
Select All Remove All	
*Name batch change:   No  Yes	

Setting	Description
Select	Select Schema Migration and Full Data Migration. In this scenario, Incremental Data Migration is not supported.
migration types	<b>Note</b> To ensure data consistency, we recommend that you do not write data to the Amazon RDS for PostgreSQL instance during data migration.
	Select one or more objects from the <b>Available</b> section and click the > icon to add
	the objects to the <b>Selected</b> section.
	⑦ Note
	• You can select columns, tables, or databases as the objects to migrate.
Select objects to migrate	<ul> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> </ul>
	<ul> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the ApsaraDB RDS instance. For more information, see Object name mapping.
	By default, if DTS fails to connect to the source or destination database, DTS retries
Specify the retry time range for	within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
connections to the source or destination database	Note When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

## 8. In the lower-right corner of the page, click **Precheck**.

? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

(?) Note We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.

Task ID/Name:	Status: Completed	View Details   Duplicate Task   Configure Monitoring and Alerting
2019-05-24 17:01:02 Created Schema Migration <b>100%</b>		2019-05-27 15:54:33 Completed Full Data Migration 100%(Migrated Rows: 1000000)
Start Pause Stop Delete		Total: 1 item(s), Per Page: 20 item(s)

12. Switch your workloads to the ApsaraDB RDS for PostgreSQL instance.

# 7.5. Migrate full data from an Amazon RDS for SQL Server instance to an ApsaraDB RDS for SQL Server instance

This topic describes how to migrate full data from an Amazon RDS for SQL Server instance to an ApsaraDB RDS for SQL Server instance by using Data Transmission Service (DTS).

### Prerequisites

- The **Public accessibility** option of the Amazon RDS for SQL Server instance is set to **Yes**. This ensures that DTS can access the Amazon RDS for SQL Server instance over the Internet.
- The database version of the Amazon RDS for SQL Server instance is 2005, 2008, 2008 R2, 2012, 2014, or 2016.
- An ApsaraDB RDS for SQL Server instance is created. For more information, see Create an ApsaraDB RDS for SQL Server instance.
- The available storage space of the ApsaraDB RDS for SQL Server instance is larger than the total size of the data in the Amazon RDS for SQL Server instance.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- Amazon RDS for SQL Server does not support incremental data migration because the sysadmin role is unavailable in Amazon RDS.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If the name of the source database is invalid, you must create a database in the ApsaraDB RDS for SQL Server instance before you configure a data migration task.

(?) Note For more information about how to create a database and the database naming conventions, see Create an account and a database for an ApsaraDB RDS instance that runs SQL Server 2012, 2016, 2017 SE, or 2019 SE.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## **Migration types**

• Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, view, trigger, synonym, SQL stored procedure, SQL function, plan guide, user-defined type, rule, default, and sequence.

• Full data migration

DTS migrates historical data of the required objects from the Amazon RDS for SQL Server instance to the ApsaraDB RDS for SQL Server instance.

⑦ Note DTS does not migrate data of the SQL\_VARIANT type.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Pricing.

## Permissions required for database accounts

Database	Schema migration	Full data migration
Amazon RDS for SQL Server	The SELECT permission	The SELECT permission
ApsaraDB RDS for SQL Server	The read and write permissions	The read and write permissions

For more information about how to create and authorize a database account, see the following topics:

- Amazon RDS for SQL Server instance: CREATE USER
- ApsaraDB RDS for SQL Server instance: Create an account for an ApsaraDB RDS for SQL Server instance.

## Process of full data migration

To prevent data migration failures caused by dependencies among objects, DTS migrates the schemas and data of the source SQL Server database in the following order:

- 1. Migrate the schemas of tables, views, synonyms, user-defined types, rules, defaults, and plan guides.
- 2. Perform full data migration.
- 3. Migrate the schemas of SQL stored procedures, SQL functions, triggers, and foreign keys.

## Before you begin

- 1. Log on to the Amazon RDS Management Console.
- 2. Go to the **Basic Information** page of the Amazon RDS for SQL Server instance.
- 3. In the **Security group rules** section, click the name of the security group corresponding to the existing inbound rule.

Security group rules (2)		C	
<b>Q</b> Filter security group rules		< 1 >	0
Security group	Туре	▼ Rule	•
	CIDR/IP - Inbound		
MARKED AND ADDRESS OF	CIDR/IP - Outbound	0.0.0.0/0	

4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

	EC2 Dashboard		Create Security Gr	oup Actions 👻					2	6 2	¢ 0
	Tags		Q Filter by tags and	attributes or search by	keyword					to 1 of 1	
	Reports		Name -	Group ID	▲ Gro	up Name	· VPC ID	•	Owner		- Desc
	INSTANCES			10000	defa	ult	-		-		defau
	Instances										
	sp Edit inbou	und r	ules							×	
	Re De Type () 3		Protocol (i) P	ort Range (i)	Source (j)			Description (i	)		
	Ca Re	F.	ТСР	1433	Custom •			dts		⊗	
•	IM. Add Rule			4	5	6					
1	AN NOTE: Any edi Bu on that rule to	ts made be drop	on existing rules will n ped for a very brief pe	esult in the edited rule riod of time until the r	e being deleted a lew rule can be c	nd a new rule create reated.	ed with the new deta	ils. This will cause t	raffic that depend		
	EL STI								Cancel	Save	
	V0 Snanshots										•
	Lifecycle Manager		Security Group:							-	
	NETWORK &		Description Inb	Outbound	Tags						
L	Security Groups		Edit 2								
	Placement Groups		Type (j)	Protoco	1	Port Range (i	Source	e (j	Description	()	
	Key Pairs	-	All traffic	All		All	100	(default)			-

### ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region. You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- You can add all of the required CIDR blocks to the inbound rule at a time.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source a	nd Destinatior	n 🔰 2.Config	ire Migration Types and Objects $ig>$	3.Map name modification	A.Precheck		
*	Task Name: S	QL Server_TO_RDS					
Source Database							
*	Instance Type:	User-Created Databa	se with Public IP Address				
* In	istance Region:	Singapore	Ŧ	Get IP Address Segment of DTS			
*[	Database Type:	SQLServer	۲				
* Hostname	e or IP Address:						
•	* Port Number:	1433					
* Data	abase Account:	dtstest		Test Connectivity @ Doccod			
Data	base Password.	**********	47	Test Connectivity Passed			
Destination Database							
*	Instance Type:	RDS Instance	Ŧ				
* In	istance Region:	Singapore	¥				
* RE	OS Instance ID:		•				
* Dat	abase Account:	dtstest					
* Datal	base Password:	•••••	4>	Test Connectivity 🔗 Passed			
					Cancel Set Whitelist and Next		
Section	Param	neter	Description				
N/A	Task I	Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.				
	Instar	се Туре	Select User-Created Dat	tabase with Public	IP Address.		
	lnstar Regio	nce n	If the instance type is set IP Address, you do not r	to <b>User-Created D</b> need to specify the ii	atabase with Public nstance region.		
	Datab	ase Type	Select SQL Server.				

## Data Migration Migrate data from a third-party cloud to Alibaba Cloud

Section	Parameter	Description
Source Dat abase	Hostname or IP Address	Enter the endpoint that is used to access the Amazon RDS for SQL server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the Amazon RDS for SQL Server instance.         Image: Transmitting the endpoint on the Basic information page of the endpoint informa
		Recommendations
	Port Number	instance. The default port number is <b>1433</b> .
	Dat abase Account	Enter the database account of the Amazon RDS for SQL Server instance. For information about permissions required for the account, see Permissions required for database accounts.
		Enter the password of the database account.
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
	Instance Type	Select RDS Instance.
	Instance Region	Select the region where the ApsaraDB RDS for SQL Server instance resides.

Section	Parameter	Description			
	RDS Instance ID	Select the ID of the ApsaraDB RDS for SQL Server instance.			
Destinatio n Database	Database Account	Enter the database account of the ApsaraDB RDS for SQL Server instance. For information about the permissions that are required for the account, see Permissions required for database accounts.			
		Enter the password of the database account.			
	Dat abase Password	<b>?</b> Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.			

### 6. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note DTS adds the CIDR blocks of DTS servers to the whitelist of the ApsaraDB RDS for SQL Server instance. This ensures that DTS servers can connect to the ApsaraDB RDS for SQL Server instance.

### 7. Select the migration types and the objects to be migrated.

igure Source and Destination 2.Configure Mig	ration Types and Objects	3.Map name modification	$\geq$	4.Precheck
<ul> <li>Migration Types: Schema Migration If Schema Migration If Schema Migration applies to short-term migration sce between Apsara Stack databases.</li> <li>For long-term data synchronization in real time, us</li> </ul>	ull Data Migration Increment narios. Typical scenarios include migra e the data synchronization feature.	al Data Migration	databases, and mig	ating data
Available       If you search globally, please expand the       If you search globally, please expand the		Selected (To edit an object name or it Edit.) Learn more.	s filter, hover over t	he object and click
Select All		Remove All		
*Name batch change:  No Yes Information: 1. Data migration only copies the data and schema in in the source database. 2. DDL operations are not supported during data mig	the source database and saves the co	py in the destination database. The process	does not affect any	data or schema
2. OD: operations are not supported during data mig	ation because this can cause migratio	n raitures.	Coursel Dava	

Setting	Description
	Select Schema Migration and Full Data Migration.
Select the migratio n types	<ul> <li>Note</li> <li>Amazon RDS for SQL Server does not support incremental data migration because the sysadmin role is unavailable in Amazon RDS.</li> <li>To ensure data consistency, we recommend that you do not write data to the Amazon RDS for SQL Server instance during data migration.</li> </ul>
	Select one or more objects from the <b>Available</b> section and click the <b>y</b> icon to move the
	objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
the source or destinati on databas	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
е	

8. In the lower-right corner of the page, click **Precheck**.

? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

Ote We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.

Task ID/Name:	Status: Completed	View Details   Duplicate Task   Configure Monitoring and Alerting
2019-05-24 17:01:02 Created Schema Migration <b>100%</b>		2019-05-27 15:54:33 Completed Full Data Migration 100%(Migrated Rows: 1000000)
Start Pause Stop Delete		Total: 1 item(s), Per Page: 20 item(s)

12. Switch your workloads to the ApsaraDB RDS for SQL Server instance.

# 7.6. Migrate data from an Amazon Aurora MySQL cluster to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from an Amazon Aurora MySQL cluster to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. You can select all of the supported migration types to ensure service continuity.

### Prerequisites

- The Public accessibility option of the Amazon Aurora MySQL cluster is set to **Yes**. The setting ensures that DTS can access the Amazon Aurora MySQL cluster over the Internet.
- An ApsaraDB RDS for MySQL instance is created. For more information, see Create an ApsaraDB RDS for MySQL instance.
- The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total size of the data in the Amazon Aurora MySQL cluster.

### Precautions

> Document Version: 20220712

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- DTS automatically creates a destination database in the ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the ApsaraDB RDS for MySQL instance before you configure the data migration task.

(?) Note For more information about the naming conventions of ApsaraDB RDS for MySQL databases and how to create a database, see Create a database on an ApsaraDB RDS for MySQL instance.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee		
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from		
ncremental data Charged. For more information, see Pricing.		information, see Pricing.		

## **Migration types**

• Schema migration

DTS migrates the schemas of the required objects to the ApsaraDB RDS for MySQL instance. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function. DTS does not support schema migration for events.

⑦ Note

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- DTS does not migrate user information. To call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

### • Full dat a migration

DTS migrates historical data of the required objects from the Amazon Aurora MySQL cluster to the ApsaraDB RDS for MySQL instance.

**?** Note During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination instance. After full data migration is complete, the tablespace of the destination instance is larger than that of the source instance.

### • Incremental data migration

After full data migration is complete, DTS retrieves binary log files from the Amazon Aurora MySQL cluster. Then, DTS synchronizes incremental data from the Amazon Aurora MySQL cluster to the ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data between MySQL databases.

## Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Amazon Aurora MySQL	The SELECT permission on the objects to be migrated	The SELECT permission on the objects to be migrated	The SELECT permission on the objects to be migrated, the REPLICATION SLAVE permission, the REPLICATION CLIENT permission, and the SHOW VIEW permission
ApsaraDB RDS for MySQL	The read and write permissions on the objects to be migrated	The read and write permissions on the objects to be migrated	The read and write permissions on the objects to be migrated

For more information about how to create and authorize a database account, see the following topics:

- Amazon Aurora MySQL cluster: Create an account for a user-created MySQL database and configure binary logging
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

## Before you begin

- 1. Log on to the Amazon Aurora console.
- 2. Go to the **Basic information** page of the Amazon Aurora MySQL cluster.
- 3. Select the node that assumes the **writer** role.
- 4. In the **Connectivity & security** section, click the name of the VPC security group that corresponds to the writer node.

<b>DB</b> identifier	•	Role 🔻	Engine 🔻	Class 🔻
		Cluster	Aurora PostgreSQL	-
0		Writer	Aurora PostgreSQL	db.r4.large
				Þ
Connectivity & security Mon	itoring Logs & events	Configurat	tion Tags	
Connectivity & security				
Endpoint & port	Networking		Security	
Endpoint	Availability zone		VPC security groups	
ap- northeast-2.rds.amazonaws.com	ap-northeast-2c			
	VPC		( active )	
Port	-		Public accessibility	
	Subnet group		Yes	

5. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

	EC2 Dashboard	Create S	Security Group	Actions 👻						∆ ⊙	ø	0
	Tags	Q, Filte	r by tags and attr	ributes or search b	y keyword				<b>0</b> K <	1 to 1 of 1		
	Reports	Na	ime 🔹 (	Group ID	<b>~</b>	Group Name	· VPC ID		• Owner		•	Desc
	Limits				(	default	10.000		-			defau
	Instances											
l	sp Edit inbound	d rules								×		
	De Type (i) 3	Protocol	(i) Port	Range (j)	Source ()			Description	0			
	Ca Re	TCP	330	6	Custom	•		dts		⊗		
-	IM. Add Rule			4	5	6						
	AN NOTE: Any edits m Bu on that rule to be d	ade on existin Iropped for a v	g rules will resu very brief period	It in the edited ru d of time until the	le being delete new rule can b	d and a new rule cre e created.	ated with the new de	etails. This will caus	e traffic that depe	nds		
	EL STI Vo								Cancel	Save		•
	Snapshots	Security	Group:	-								1
	NETWORK & SECURITY	Descrip	tion	nd Outbound	d Tags							
L	Security Groups	Edit	2									
	Placement Groups	Туре	()	Protoco	ol (i)	Port Range	(j) Sou	rce (i)	Descriptio	on (j		
	Key Pairs	All tra	ffic	All		All		(default)				-

### ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region.
   You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- You can add all of the required CIDR blocks to the inbound rule at a time.
- 6. Log on to the Amazon Aurora MySQL database and specify the number of hours to retain binary log files. Skip this step if you do not need to perform incremental data migration.

call mysql.rds\_set\_configuration('binlog retention hours', 24);

### ? Note

- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of the Amazon Aurora MySQL cluster must be enabled and the value of the binlog\_format parameter must be set to row. If the MySQL version is 5.6 or later, the value of the binlog\_row\_image parameter must be set to full.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

## Data Migration Migrate data from a

Cancel

Set Whitelist and Next

third-party	/cloudto	Alibaba	Cloud
-------------	----------	---------	-------

				- M		t Durch a de			
1.Configure Source and Destination	2.Configure Migration Types and Objects	>	3.Map name modific	atton		4.Precneck			
* Task Name: M	Task Name: MySQL_TO_MySQL								
Source Database									
* Instance Type:	User-Created Database with Public IP Address	Ŧ							
* Instance Region:	Singapore	۳	Get IP Address Segment of	f DTS					
* Database Type:	MySQL	•							
* Hostname or IP Address:									
* Port Number:	3306								
* Database Account:	dtstest								
* Database Password:	•••••	4>	Test Connectivity	⊘ Passed					
Destination Database									
<ul> <li>Instance Type:</li> </ul>	RDS Instance	٣							
* Instance Region:	Singapore	•							
* RDS Instance ID:		•							
* Database Account:	dtstest								
* Database Password:	•••••	<b>∮</b> >	Test Connectivity	⊘ Passed					

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select User-Created Database with Public IP Address.
	lnstance Region	If the instance type is set to <b>User-Created Database with Public</b> <b>IP Address</b> , you do not need to specify the <b>instance region</b> .
	Database Type	Select MySQL.

Section	Parameter	Description					
		Enter the endpoint that is used to access the Amazon Aurora MySQL cluster.					
		Note You can obtain the endpoint on the Basic information page of the Amazon Aurora MySQL cluster.					
Source Dat <i>a</i> base	Hostname or IP Address	DB identifier       Role       Engine       Class         Cluster       Aurora PostgreSQL       -         Writer       Aurora PostgreSQL       dbr4.large         Mriter       Aurora PostgreSQL       dbr4.large         Connectivity & security       Monitoring       Logs & events       Configuration       Tags					
		Connectivity & security					
		Endpoint & port     Networking     Security       Endpoint     Availability zone     VPC security groups       ap-northeast-2c     Image: Construct of the security groups       Port     VPC     (active )       Port     Public accessibility       Subnet group     Yes					
	Port Number	Enter the service port number of the Amazon Aurora MySQL cluster. The default port number is <b>3306</b> .					
	Dat abase Account	Enter the database account of the Amazon Aurora MySQL cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.					
	Dat abase Password	Enter the password of the database account. <b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.					
	Instance Type	Select RDS Instance.					
	Instance	Select the region where the ApsaraDB RDS for MySQL instance resides.					
	RDS Instance ID	Select the ID of the ApsaraDB RDS for MySQL instance.					

Section	Parameter	Description						
	Database Account	Enter the database account of the ApsaraDB RDS for MySQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.						
		Enter the password of the database account.						
Destinatio n Database	Dat abase Password	⑦ Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.						
	Encryption	Select Non-encrypted or SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption on an ApsaraDB RDS for MySQL instance <b>Note</b> The Encryption parameter is available only for regions in mainland China and the China (Hong Kong) region.						
	Database Password	message appears, click Check next to Failed. Then, modify the information based on the check results. Select Non-encrypted or SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the RDS instant before you configure the data migration task. For more information, see Configure SSL encryption on an ApsaraDB RDS for MySQL instance ONCE The Encryption parameter is available only for regions in mainland China and the China (Hong Kong) region.						

6. In the lower-right corner of the page, click Set Whitelist and Next.

### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

### Data Migration Migrate data from a

Dat a Transmission Service

1.Config	ure Source and	2.Configure Migration Ty	/pes and	3.Advanced Settings	$\rightarrow$	4.Precheck	
<ul> <li>Migration 1</li> <li>triggers, For n</li> </ul>	Types: 🗹 Schema Mig nore information, see R	gration 🔽 Full Data Migration	✓ Incremental [	Data Migration Note: In	cremental data migi	ration does not support	
Note: do no cleans up ti	ot clean up the increme he log too early, the Dī	ntal data log generated by the sour	e database after the	DTS task is started when th	e DTS full task is ru	nning. If the source datab	ase
Data migral between Ap For long-ter	tion applies to short-te osara Stack databases, rm data synchronizatio	rm migration scenarios. Typical scen n in real time, use the data synchror	arios include migratin iization feature.	g data to the doud, scaling	and sharding datab	ases, and migrating data	
Available				Selected (To edit an obj	ect name or its filter	r, hover over the object ar	nd clic
Expand the	e tree before you perfo	rm a gloi 🔰 🔍				Q	
🗆 📑 dtsl	testdata			🚍 dtstestdata (2	Ohiects)		
	ables /iews			customer	.00,000,		
			>	order			
			`				
Select All				Remove All			
*Rename Dat	abases and Tables:	Do Not Change Databas	e and Table Names	Change Database and	d Table Names		
* Retry Time f	for Failed Connection	720 Minute	5 🕜				
*Source table want to copy t	DMS_ ONLINE_ Do yo he temporary table to	u 🔾 Yes 💿 No 🕐					
the target dat	abase during DDL:						
Information: 1. Data migrat	: ion only copies the dat	a and schema in the source databas	e and saves the copy	in the destination database	. The process does	not affect any data or sch	ema
in the source o 2. Do not do D	database. DL operation during st	ructure and full migration, otherwise	e the task may fail				
				Cancel	Previous	Save Preche	sck
etting	Descriptior	1					
	<ul> <li>To perfo</li> <li>Migration</li> </ul>	orm only full data mi on.	gration, sele	ct Schema Mig	r <b>ation</b> and	Full Data	
elect e	<ul> <li>To ensu</li> <li>Data M</li> </ul>	re service continuity o igration, and Incren	during data r nental Data	nigration, select Migration.	Schema M	ligration, Full	
igratio	<b>Note</b> If <b>Incremental Data Migration</b> is not selected, we recommend that you do not write data to the Amazon Aurora MySQL cluster during data migration. This						

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the <b>&gt;</b> icon to move the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
the source or destinati on databas e	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
е	

Setting	Description
Specify whether to copy tempora ry tables to the destinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.
on databas e when	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.
perform s online ומס	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
operatio ns on the source table	<b>Note</b> If you select No, the tables in the destination database may be locked.

### 8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the ApsaraDB RDS for MySQL instance.

# 7.7. Migrate data from an Amazon Aurora MySQL cluster to a PolarDB for MySQL cluster

This topic describes how to migrate data from an Amazon Aurora MySQL cluster to a PolarDB for MySQL cluster by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. You can select all of the supported migration types to ensure service continuity.

## Prerequisites

- The Public accessibility option of the Amazon Aurora MySQL cluster is set to **Yes**. The setting ensures that DTS can access the Amazon Aurora MySQL cluster over the Internet.
- A PolarDB for MySQL cluster is created. For more information, see Create a PolarDB for MySQL cluster.
- The available storage space of the PolarDB for MySQL cluster is larger than the total size of the data in the Amazon Aurora MySQL cluster.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.

- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the FLOAT or DOUBLE data type. If you do not specify a precision, DTS sets the precision for the FLOAT data type to 38 digits and the precision for the DOUBLE data type to 308 digits. You must check whether the precision settings meet your business requirements.
- If the name of the source database is invalid, you must create a database in the PolarDB for MySQL cluster before you configure a data migration task.

**Note** For more information about how to create a database and the database naming conventions, see **Create a database on an ApsaraDB RDS for MySQL instance**.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.	

## Migration types

### • Schema migration

DTS migrates the schemas of the required objects to the PolarDB for MySQL cluster. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function. DTS does not support schema migration for events.

### ? Note

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- DTS does not migrate user information. To call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

### • Full data migration

DTS migrates historical data of the required objects from the Amazon Aurora MySQL cluster to the PolarDB for MySQL cluster.

**?** Note During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination instance. After full data migration is complete, the tablespace of the destination instance is larger than that of the source instance.

### • Incremental data migration

After full data migration is complete, DTS retrieves binary log files from the Amazon Aurora MySQL cluster. Then, DTS synchronizes incremental data from the Amazon Aurora MySQL cluster to the PolarDB for MySQL cluster. Incremental data migration allows you to ensure service continuity when you migrate data between MySQL databases.

## Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Amazon Aurora MySQL	The SELECT permission on the objects to be migrated	The SELECT permission on the objects to be migrated	The SELECT permission on the objects to be migrated, the REPLICATION SLAVE permission, the REPLICATION CLIENT permission, and the SHOW VIEW permission
PolarDB for MySQL	The read and write permissions on the objects to be migrated	The read and write permissions on the objects to be migrated	The read and write permissions on the objects to be migrated

For more information about how to create and authorize a database account, see the following topics:

- Amazon Aurora MySQL cluster: Create an account for a user-created MySQL database and configure binary logging
- PolarDB for MySQL cluster: Create a database account.

## Before you begin

- 1. Log on to the Amazon Aurora console.
- 2. Go to the **Basic information** page of the Amazon Aurora MySQL cluster.
- 3. Select the node that assumes the **writer** role.
- 4. In the **Connectivity & security** section, click the name of the VPC security group that corresponds to the writer node.

DB identifier	▲ Role	▼ Engine ▼ Class ▼				
	Cluste	er Aurora PostgreSQL -				
0	Write	r Aurora PostgreSQL db.r4.large				
Connectivity & security Monitoring Logs & events Configuration Tags Connectivity & security						
Endpoint & port	Networking	Security				
Endpoint	Availability zone	VPC security groups				
ap- northeast-2.rds.amazonaws.com	ap-northeast-2c					
	VPC	( active )				
Port	10-00-004	Public accessibility				
	Subnet group	Yes				

5. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

	EC2 D	ashboard		Create Security (	Group Actions 👻						4 0	ø	0
	Tags			Q, Filter by tags a	nd attributes or search t	oy keyword				0 K <	1 to 1 of	1 >	
	Report	ts		Name	• Group ID	-	Group Name	- VPC ID	•	Owner		•	Desc
	Limits				10000		default	10.000					defau
	Instand	ces											
l	La Sp	Edit inbound	d ru	iles							;	×	
	De	Type (i) 3		Protocol (i)	Port Range (j)	Source ()			Description (i)				
	Ca Re	Custom TCP F •		TCP	3306	Custom	<b>v</b>		dts		⊗		
-	IM. AN Bu	Add Rule NOTE: Any edits m on that rule to be d	iade fropp	on existing rules wil bed for a very brief p	4 Il result in the edited ru period of time until the	5 ule being delete new rule can l	ed and a new rule cl be created.	6 reated with the new detai	ls. This will cause tr	affic that depe	nds		
-	EL STO Vo									Cancel	Save	]	Þ
	Snaps	hots		Security Group:	1						_		
•	NETWO	DRK &		Description	nbound Outbour	id Tags							
l	Secur	rity Groups		Edit 2									
	Placen	nent Groups		Туре (і)	Protoc	ol (j	Port Range	e (i) Source	e (i)	Descriptio	on (j		
	Key Pa	airs 🗸		All traffic	All		All	100	(default)				-

### ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region.
   You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- You can add all of the required CIDR blocks to the inbound rule at a time.
- 6. Log on to the Amazon Aurora MySQL database and specify the number of hours to retain binary log files. Skip this step if you do not need to perform incremental data migration.

call mysql.rds\_set\_configuration('binlog retention hours', 24);

### ? Note

- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of the Amazon Aurora MySQL cluster must be enabled and the value of the binlog\_format parameter must be set to row. If the MySQL version is 5.6 or later, the value of the binlog\_row\_image parameter must be set to full.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source a	nd Destination	2.Configu	ire Migration Types and Objects 🔪	3.Map name modification	4.Precheck		
	[						
	Task Name: M	VSQL_TO_POLARDB					
Source Database							
•	Instance Type:	User-Created Databa	ase with Public IP Address				
* In	stance Region:	Singapore	Ţ	Get IP Address Segment of DTS			
*[	Database Type:	MySQL	Ŧ				
* Hostname	or IP Address:						
	Port Number:	3306					
* Data	abase Account:	dtstest					
* Datab	base Password:	•••••	Ф	Test Connectivity 🔗 Passed			
Destination Database							
*	Instance Type:	PolarDB	T				
* In	stance Region:	Singapore	•				
* PolarD	B Instance ID:		-				
* Data	abase Account:	dtstest					
* Datab	base Password:	•••••	ф)	Test Connectivity 🔗 Passed			
					Cancel Set Whitelist and Next		
Section	Param	eter	Description				
N/A	Task N	Jame	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.				
	Instan	се Туре	Select User-Created Database with Public IP Address.				
	Instance Region		If the instance type is set to <b>User-Created Database with Public</b> <b>IP Address</b> , you do not need to specify the <b>instance region</b> .				
	Database Type		Select MySQL.				

Source Database         Enter the endpoint that is used to access the Amazon Aurora MySQL cluster.           Source Database         Note You can obtain the endpoint on the Basic information page of the Amazon Aurora MySQL cluster.           Image: the database of the Amazon Aurora MySQL cluster.         Image: the database of the Amazon Aurora MySQL cluster.           Image: the database of the Amazon Aurora MySQL cluster.         Image: the database of the Amazon Aurora MySQL cluster.           Port Number         Enter the service port number of the Amazon Aurora MySQL cluster. For formation about the permissions that are required for the account.           Database Password         Enter the password of the database account.           Image: the service port number of the Amazon Aurora MySQL cluster. For information about the permissions that are required for the account. See Permissions required for database account.           Image: the password of the database account.         Image: the service parameters are valid, the Passed message appears, if the Failed parameters are valid, the Passed message appears. (If the specified parameters are valid, the Passed message appears.) (If the specified parameters are valid, the Passed message appears.) (If the specified parameters are valid, the Passed message appears.) (If the Failed message appears.) (If Check next to Failed. Modify the source database parameters based on the check results.           Imatance Region         Select the region where the PolarDB for MySQL cluster resides.           PolarDB Instance ID         Select the ID of the PolarDB for MySQL cluster.	Section	Parameter	Description					
Source Database         Hostname or IP Address         Information page of the Amazon Aurora MySQL cluster.           Port Number         Enter the service port number of the Amazon Aurora MySQL cluster. The default port number is 3306.           Port Number         Enter the service port number of the Amazon Aurora MySQL cluster. The default port number is 3306.           Database Account         Enter the service port number of the Amazon Aurora MySQL cluster. The default port number is 3306.           Database Account         Enter the database account of the Amazon Aurora MySQL cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.           Database Account         Enter the password of the database account.           Image: Database Password         Image: Select PolarDB.           Instance Region         Select the region where the PolarDB for MySQL cluster.           PolarDB Instance ID         Select the ID of the PolarDB for MySQL cluster.			Enter the endpoint that is used to access the Amazon Aurora MySQL cluster.           ⑦ Note         You can obtain the endpoint on the Basic					
Port NumberEnter the service port number of the Amazon Aurora MySQL cluster. The default port number is 3306.Database AccountEnter the database account of the Amazon Aurora MySQL cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.Database PasswordEnter the password of the database account. (In Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the source database parameters based on the check results.Instance TypeSelect PolarDB.PolarDB Instance IDSelect the region where the PolarDB for MySQL cluster resides.	Source Database	Hostname or IP Address	Information page of the Amazon Aurora MySQL cluster.					
Database AccountEnter the database account of the Amazon Aurora MySQL cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.Database AccountEnter the password of the database account.Database PasswordEnter the password of the database account.Instance TypeSelect PolarDB.Instance RegionSelect the region where the PolarDB for MySQL cluster resides.PolarDB Instance IDSelect the ID of the PolarDB for MySQL cluster.		Port Number	Enter the service port number of the Amazon Aurora MySQL cluster. The default port number is <b>3306</b> .					
Database PasswordEnter the password of the database account.Database PasswordInstance TypeInstance RegionSelect PolarDB.PolarDB Instance IDSelect the region where the PolarDB for MySQL cluster resides.		Dat abase Account	Enter the database account of the Amazon Aurora MySQL cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.					
Instance TypeSelect PolarDB.Instance RegionSelect the region where the PolarDB for MySQL cluster resides.PolarDB Instance IDSelect the ID of the PolarDB for MySQL cluster.		Dat abase Password	Enter the password of the database account. <b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.					
Instance RegionSelect the region where the PolarDB for MySQL cluster resides.PolarDB Instance IDSelect the ID of the PolarDB for MySQL cluster.		Instance Type	Select <b>PolarDB</b> .					
PolarDB Instance ID Select the ID of the PolarDB for MySQL cluster.		Instance Region	Select the region where the PolarDB for MySQL cluster resides.					
		PolarDB Instance ID	Select the ID of the PolarDB for MySQL cluster.					

Section	Parameter	Description			
Destinatio n Database	Dat abase Account	Enter the database account of the PolarDB for MySQL cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.			
		Enter the password of the database account.			
	Database Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.			

### 6. In the lower-right corner of the page, click **Set Whitelist and Next**.

**Note** In this step, DTS adds the CIDR blocks of DTS servers to the whitelist of the PolarDB for MySQL cluster. This ensures that DTS servers can connect to the cluster.

### 7. Select the migration types and the objects to be migrated.

### Dat a Transmission Service

## Data Migration Migrate data from a third-party cloud to Alibaba Cloud

1.Configu	re Source and	2.	Configure Migration Ty	pes and	3.Advance	d Settings	$\geq$	4.Precheck
* Migration Ty triggers, For mo	ypes: 🗹 Schema M ore information, see	ligration Reference	Full Data Migration	✓ Incremental [	Data Migration	Note: Increm	ental data migral	tion does not support
Note: do not cleans up the	clean up the increme e log too early, the [	nental data l DTS increme	og generated by the souro intal task may fail	e database after the	DTS task is start	ted when the DT:	5 full task is runr	ning. If the source database
Data migrati between Aps For long-terr	on applies to short-t sara Stack databases n data synchronizati	erm migrati 5. on in real tir	on scenarios. Typical scena me, use the data synchroni	inos include migratir	ig data to the do	ud, scaling and s	sharding databas	es, and migrating data
Available					Selected (To Edit.) Learn	edit an object na	ame or its filter,	hover over the object and cli
Expand the	tree before you per	form a glol	Q					Q
🗆 💼 dtste	estdata ablos				🕋 dtste	stdata (20bj	ects)	
E 🔤 Vi	iews				🔳 cust	omer		
				>	🔳 orde	r		
				-				
				<				
Select All					Pamous All			
Rename Data	bases and Tables:	۲	Do Not Change Database	and Table Names	Change E	atabase and Tab	ble Names	
Retry Time fo	or Failed Connection	[	720 Minutes	0				
Source table D	DMS_ONLINE_Do y	ou O	Yes 💿 No 🕜					
vant to copy th he target datal	he temporary table to base during DDL:	0						
oformation:	-							
L. Data migratio	on only copies the d	ata and sche	ema in the source database	e and saves the copy	in the destination	on database. The	process does no	ot affect any data or schema
. Do not do Di	DL operation during	structure an	d full migration, otherwise	the task may fail				
						Cancel	Previous	Save Precheck
tting	Descriptio	n						
	• To perf	orm or	nly full data mig	gration, sele	ct Schem	a Migrat	ion and F	ull Data
	Migrat	ion.						
lect	<ul> <li>To ensi</li> <li>Data M</li> </ul>	ure serv ligrati	vice continuity d on, and Increm	uring data r ental Data	nigration, Migratio	select <b>Sc</b> on.	hema Mi	gration, Full
migratio								

ensures data consistency between the source and destination databases.

## Data Migration-Migrate data from a third-party cloud to Alibaba Cloud

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. Otherwise, the data migration task fails.
the source or destinati on databas	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

Setting	Description		
Specify whether to copy tempora ry tables to the destinati on databas e when DMS perform s online DDL operatio ns on the source table	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.		
	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.		
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.		
	<b>Note</b> If you select No, the tables in the destination database may be locked.		

### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the PolarDB for MySQL cluster.

# 7.8. Migrate full data from an Amazon Aurora PostgreSQL instance to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to migrate full data from an Amazon Aurora PostgreSQL instance to an ApsaraDB RDS for PostgreSQL instance by using Data Transmission Service (DTS).

### Prerequisites

- An Amazon Aurora PostgreSQL instance is created and the database version is 9.4, 9.5, 9.6, or 10.0.
- The **Public accessibility** option of the Amazon Aurora PostgreSQL instance is set to **Yes**. This ensures that DTS can access the Amazon Aurora PostgreSQL instance over the Internet.
- An ApsaraDB RDS for Post greSQL instance is created. For more information, see Create an ApsaraDB RDS for Post greSQL instance.

Note The database engine version of the ApsaraDB RDS for PostgreSQL instance is 9.4 or 10.0.

• The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total size of the data in the Amazon Aurora PostgreSQL instance.

### Precautions

• DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no

primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.

- You cannot use DTS to migrate incremental data from an Amazon Aurora PostgreSQL instance to an ApsaraDB RDS for PostgreSQL instance. Before you start the data migration task, you must stop the services that run on the Amazon Aurora PostgreSQL instance. To ensure data consistency, we recommend that you do not write data to the Amazon Aurora PostgreSQL instance during data migration.
- Each data migration task can migrate data from only a single database. To migrate data from multiple databases, you must create a data migration task for each database.
- Functions that are written in the C programming language cannot be migrated.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If a data migration task fails, DTS automatically resumes the task. Therefore, before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance overwrites the data in the destination instance after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Pricing.

## **Migration types**

• Schema migration

DTS migrates the schemas of required objects to the ApsaraDB RDS for PostgreSQL instance. DTS supports the following types of objects for schema migration: table, trigger, view, sequence, function, user-defined type, rule, domain, operation, and aggregate.

• Full data migration

DTS migrates the historical data of required objects from the Amazon RDS for PostgreSQL instance to the ApsaraDB RDS for PostgreSQL instance.

## Permissions required for database accounts

Database	Schema migration	Full data migration
Amazon Aurora PostgreSQL	The USAGE permission on pg_catalog	The SELECT permission on the objects to be migrated
ApsaraDB RDS for PostgreSQL	The CREATE and USAGE permissions on the objects to be migrated	The permissions of the schema owner

## Process of full data migration
To prevent data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

- 1. Migrate the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates.
- 2. Perform full data migration.
- 3. Migrate the schemas of triggers and foreign keys.

#### Preparation 1: Edit the inbound rule of the Amazon Aurora PostgreSQL instance

- 1. Log on to the Amazon Aurora console.
- 2. Go to the Basic Information page of the Amazon Aurora PostgreSQL instance.
- 3. Select the node that assumes the writer role.
- 4. In the **Connectivity & security** section, click the name of the VPC security group that corresponds to the writer node.

DB identifier	▲ I	Role 🔻	Engine 🔻	Class
	(	Cluster	Aurora PostgreSQL	-
•		Writer	Aurora PostgreSQL	db.r4.large
Connectivity & security Mon	itoring Logs & events	Configurat	tion Tags	
Connectivity & security				
<b>Connectivity &amp; security</b> Endpoint & port	Networking		Security	
Connectivity & security Endpoint & port	Networking Availability zone		Security VPC security groups	
Connectivity & security Endpoint & port Endpoint	Networking Availability zone ap-northeast-2c		Security VPC security groups	
Connectivity & security Endpoint & port Endpoint int int int int int int int int int	Networking Availability zone ap-northeast-2c		Security VPC security groups	
Connectivity & security Endpoint & port Endpoint Endpoint ap- northeast-2.rds.amazonaws.com	Networking Availability zone ap-northeast-2c VPC		Security VPC security groups ( active )	
Connectivity & security Endpoint & port Endpoint Endpoint ap- northeast-2.rds.amazonaws.com Port	Networking Availability zone ap-northeast-2c VPC		Security VPC security groups ( active )	

5. On the Security Groups page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit to add the CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

EC2 Dashboard	Actions V				0	¢ 0
Tags Q. Filter by tags and attribut	es or search by keyword				:o 1 of 1	
Reports Limits	up ID 🔺	Group Name	✓ VPC ID	- Owner		• Desc
INSTANCES Instances		detault	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1			derau
sp Edit inbound rules					×	
De Type () 3 Protocol () Port Ra	nge (j) Source (j)		Descriptio	n (j)		
Re Custom TCP F TCP 5432	Custom		dts		$\otimes$	
IM         Add Rule           AN         NOTE: Any edits made on existing rules will result in Bu           Bu         on that rule to be dropped for a very brief period of	the edited rule being delete time until the new rule can be	d and a new rule created w e created.	ith the new details. This will ca	use traffic that depends	7	
EL ST Vo				Cancel	Save	•
Snapshots Security Group:		5.8.8				
NETWORK & Description Inbound	Outbound Tags					
Security Groups Edit 2						
Placement Groups Type (i)	Protocol (j)	Port Range (i)	Source (j)	Description	()	
Key Pairs All traffic	All	All	(defaul	t)		-

#### ? Note

- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region. You need to add only the CIDR blocks of DTS servers that reside in the China (Hangzhou) region.
- You can add all of the required CIDR blocks to the inbound rule at a time.

# Preparation 2: Create a database and schema in the destination RDS instance

Create a database and schema in the destination RDS instance based on the database and schema information of the objects to be migrated. The schema name of the source and destination databases must be the same. For more information, see Create a database on an ApsaraDB RDS for PostgreSQL instance and Appendix: User and schema management.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

### Data Migration Migrate data from a third-party cloud to Alibaba Cloud

1.Set Source and Desti	ination Databases	2.Mig	ration Types and Tasks		3.Precheck	
* Task Name:	(george					
Source Database						
* Instance Type:	User-Created Database with Public IP Address	•				
* Source Instance Region:	Singapore	٣	Obtain IP Address Segment of DTS			
* Database Type:	PostgreSQL					
* Hostname or IP Address:						
* Port Number:	5432					
* Database Name:	dtstestdata					
* Database Account:						
* Database Password:	•••••	<⊅	Test Connection			
Destination Database						
* Instance Type:	RDS Instance	٣				
* Source Instance Region:	China (Hangzhou)	•				
* RDS Instance ID:		•				
* Database Name:	dtstestdata					
* Database Account:						
* Database Password:	•••••	<⊅	Test Connection			
				Cancel	Assess Migration	Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select User-Created Database with Public IP Address.
	Instance Region	Select the region where the source instance resides. If the instance type is set to <b>User-Created Database with Public IP Address</b> , you do not need to specify the instance region.
Database Type		Select PostgreSQL.

Section	Parameter	Description				
		Enter the endpoint that is used to access the Amazon Aurora PostgreSQL instance.				
Source		<ul> <li>Note You can find the endpoint on the Basic</li> <li>Information page of the Amazon Aurora PostgreSQL instance.</li> </ul>				
Database		Related Q Filter databases				
	Hostname or IP Address	Db identifier     ▲     Role ▼     Engine     ▼     Class ▼     Status ▼     CPU       ●     □				
		Connectivity & security Monitoring Logs & events Configuration Maintenance & backups Tags				
		Endpoints (2)       Edit     Delete     Create custom endpoint       Q. Filter endpoint     < 1 > 6				
		Endpoint name     Status     V     Type     Port       database-1.cluster-ro-     us-west-1.rds.amazonaws.com     Available     Reader     5432				
		database-1.cluster: us-west-1.rds.amazonaws.com Available writer 5432				
	Port Number	Enter the service port number of the Amazon Aurora PostgreSQL instance. The default port number is <b>5432</b> .				
	Database Name	Enter the name of the source database in the Amazon Aurora PostgreSQL instance.				
	Dat abase Account	Enter the database account of the Amazon Aurora PostgreSQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.				
		Enter the password of the database account.				
	Dat abase Password	Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the source database parameters based on the check results.				
Instance Type Select RDS Instan		Select RDS Instance.				
	Instance Region	Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.				
RDS Instance Select the ID of the des ID instance.		Select the ID of the destination ApsaraDB RDS for PostgreSQL instance.				

Section	Parameter	Description	
		Enter the name of the destination database in the ApsaraDB RDS for PostgreSQL instance. The name can be different from the name of the source database in the Amazon Aurora PostgreSQL instance.	
Destination Database	Database Name	Note Before you configure the data migration task, you must create a database and schema in the ApsaraDB RDS for PostgreSQL instance. For more information, see Preparation 2: Create a database and schema in the destination RDS instance.	
	Dat abase Account	Enter the database account of the destination ApsaraDB RDS for PostgreSQL instance. For information about the permissions that an required for the account, see Permissions required for database accounts.	
		Enter the password of the database account.	
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.	

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

Pilgrauon rypes. V Schema Pilgrauon V Fun Data Pilgrauon		
Data migration applies to short-term migration scenarios. Typical sce between Apsara Stack databases.	enarios include migratir	ig data to the cloud, scaling and sharding databases, and migrating data
For long-term data synchronization in real time, use the data synchr	onization feature.	
Available		<b>Selected</b> (To edit an object name or its filter, house aver the object and slick
Available		Edit.) Learn more.
If you search globally, please expand the $\hfill \hfill $		Q
🕀 💼 mytest		MYTECT Source Database
		TABLE3
	>	
	<	
Select All		Remove All
*Name batch change:  No Ves		Renove An
Information: 1. Data migration only copies the data and schema in the source datab	ase and saves the copy	in the destination database. The process does not affect any data or schema
in the source database.		

Setting	Description
Select the	Select <b>Schema Migration</b> and <b>Full Data Migration</b> . In this data migration task, <b>Incremental Data Migration</b> is not supported.
migration types	<b>Note</b> To ensure data consistency, we recommend that you do not write data to the Amazon Aurora PostgreSQL instance during data migration.
	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
Select the objects to be migrated	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
connections to the source or destination database	Note When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

#### 8. In the lower-right corner of the page, click **Precheck**.

? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

Ote We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.

Task ID/Name:	Status: Completed	View Details   Duplicate Task   Configure Monitoring and Alerting
2019-05-24 17:01:02 Created Schema Migration <b>100%</b>		2019-05-27 15:54:33 Completed Full Data Migration 100%(Migrated Rows: 1000000)
Start Pause Stop Delete		Total: 1 item(s), Per Page: 20 item(s)

12. Switch your workloads to the ApsaraDB RDS for PostgreSQL instance.

## 7.9. Migrate data from a MongoDB Atlas database to an ApsaraDB for MongoDB instance

This topic describes how to use Data Transmission Service (DTS) to migrate incremental data from a MongoDB Atlas database to an ApsaraDB for MongoDB instance. DTS supports full data migration and incremental data migration. The data migration does not interrupt your services if you select the two migration types.

#### Prerequisites

The storage capacity of the destination ApsaraDB for MongoDB instance is larger than the total size of the data in the MongoDB Atlas database.

#### Precautions

• DTS consumes the resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If you migrate a large volume of data or if the server specifications do not meet your requirements, database services may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and

destination databases. We recommend that you migrate data during off-peak hours.

- DTS cannot migrate data from the admin or local database.
- The config database is an internal database. We recommend that you do not migrate this database.
- If the source and destination MongoDB databases use different versions or storage engines, make sure that your applications can run on both databases. For more information about the versions and storage engines that are supported by ApsaraDB for MongoDB, see MongoDB versions and storage engines.

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba
Incremental data migration	Charged. For more information, see Pricing.	Cloud over the Internet. For more information, see Pricing.

#### Migration types

Migration type	Description
	DTS migrates the historical data of required objects from the source MongoDB database to the destination MongoDB database.
Full data migration	<b>Note</b> The following types of objects are supported: database, collection, and index.
	After full data migration is complete, DTS synchronizes incremental data from the source MongoDB database to the destination MongoDB database.
Incremental data migration	<ul> <li>Note</li> <li>The create and delete operations that are performed on databases, collections, and indexes can be synchronized.</li> <li>The create, delete, and update operations that are performed on documents can be synchronized.</li> </ul>

#### Permissions required for database accounts

Database	Full data migration	Incremental data migration
Source MongoDB Atlas database	Read permissions on the source database and the permission to perform the listDatabases operation	<ul> <li>Read permissions on the source, admin, and local databases</li> <li>The permission to perform the listDatabases operation</li> </ul>

Database	Full data migration	Incremental data migration
Destination ApsaraDB for MongoDB instance	Read and write permissions on the destination database	Read and write permissions on the destination database

For more information about how to create a database account and grant permissions to the account, see the following topics:

- MongoDB Atlas databases: db.createUser()
- ApsaraDB for MongoDB instances: Manage user permissions on MongoDB databases

#### Preparations

- 1. Log on to the MongoDB Atlas console.
- 2. In the left-side navigation pane, click **Network Access**. On the page that appears, click **ADD IP ADDRESS**.

CONTEXT	> PROJECT 0		
Project 0 👻	Network Access		
ATLAS	IP Whitelist Peering		
Data Lake BETA			+ ADD IP ADDRESS
SECURITY Database Access	You will only be able to connect to your cluster	rom the following list of IP Addresses:	2
Network Access	1 Address Com	nent Status	Actions
Advanced		Active	CEDIT DELETE
PROJECT			
Access Management			

3. In the dialog box that appears, click ALLOW ACCESS FROM ANYWHERE and click Confirm.

Add Whitelist Entry		
Add a whitelist entry using eith ADD CURRENT IP ADDRESS	ALLOW ACCESS FROM ANYWHERE	
Whitelist Entry:	0.0.0/0	
Comment:	Optional comment describing this entry	
Save as temporary whitelist	Cancel Confirm	

Onte This step allows all IP addresses to access the MongoDB Atlas database. Delete this rule after data migration is complete.

#### Procedure

- 1. Purchase a data migration instance. For more information, see Purchase a data migration instance.
- 2. Log on to the DTS console.
- 3. In the left-side navigation pane, click **Data Migration**.
- 4. On the upper part of the **Migration Tasks** page, select the region where the ApsaraDB for MongoDB instance resides.
- 5. Find the purchased data migration instance and click **Configure Migration Task**.
- 6. Configure the source and destination databases.

Create Migration Task & Back							
1.Source en	ndpoint a	and target endpoint		2.Migration class and list		>	3.Pre-check
* Task Na	ame: a	tlas_to_aliyun					
Source Database							
* Instance	e Type:	On-premises Databases	٣				
* Instance R	Region:		•	Get DTS IP			
* Database E	Engine:	MongoDB	•				
* Host Name or IP Ad	ddress:	cluster0-shard-	mongodb.net				
	* Port:	27017					
Database	Name:	admin		Account verification database			
Database ac	count:	atlastest					
Database Pass	sword:	*****	4>	Test the Connection	⊘ Passed		
Target Database							
* Instance	e Type:	MongoDB Instance	v				
* Instance R	Region:	- Name of the State of the Stat	•				
* MongoDB Instan	nce ID:		-				
* Database I	Name:	admin		Account verification database			
* Database ac	count:	root					
* Database Pas	sword:	•••••	<b>4</b> >	Test the Connection	⊘ Passed		
Cancel Authorize Whitelist and Enter into Next Step							
Section	Pá	arameter	Description	1			
			The task na	ame that DTS a	utomatical	ly generates. V	We recommend

N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.
Instance Type		Select User-Created Database with Public IP Address.
	Instance Region	If you select <b>User-Created Database with Public IP Address</b> as the access method of the source database, you do not need to specify the <b>Instance Region</b> parameter.
	Database Type	Select MongoDB.

Section	Parameter	Description			
Hos IP A Source Database	Hostname or IP Address	The endpoint of the PRIMARY node in the MongoDB Atlas database. You can obtain the endpoint in the MongoDB Atlas console, as shown in the following figure.			
	Port Number	The service port number of the MongoDB Atlas database. Default value: <b>27017</b> .			
	Database Name	The name of the authentication database. The database account is created in this database.			
	Dat abase Account	The account of the MongoDB Atlas database. For information about the permissions that are required for the account, see Permissions required for database accounts.			
		The password of the database account.			
	Dat abase Password	<b>Note</b> After you specify the information about the source database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is correct. If the information is correct, the <b>Passed</b> message is displayed. If the information is incorrect, the <b>Failed</b> message is displayed and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.			
	Encryption	Specifies whether to encrypt the connection. In this example, <b>SSL</b> -encrypted is selected.			
	Instance Type	Select MongoDB Instance.			
	Instance Region	The region where the destination instance resides.			
	MongoDB Instance ID	The ID of the destination instance.			

Section	Parameter	Description	
Database Destinatio Name n Database		The name of the authentication database. The database account is created in this database.	
	<b>Note</b> If you want to use the root account, enter admin in the Database Name field.		
	Database	The database account of the ApsaraDB for MongoDB instance. For	
Account	information about the permissions that are required for the account, see Permissions required for database accounts.		
		The password of the database account.	
	Database Password	The password of the database account. <b>Note</b> After you specify the information about the destination database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is correct. If the information is correct, the <b>Passed</b> message is displayed. If the information is incorrect, the <b>Failed</b> message is displayed and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.	

#### 7. In the lower-right corner of the page, click Set Whitelist and Next.

#### Q Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 8. Select the migration type and the objects to migrate.

Setting	Description
Select the migration types	<ul> <li>To perform only full data migration, select Full Data Migration.</li> <li>To migrate data with minimal downtime, select both Full Data Migration and Incremental Data Migration.</li> <li>Note If Incremental Data Migration is not selected, we recommend that you do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.</li> </ul>
Select the objects to be migrated	<ul> <li>Select one or more objects from the Source Objects section and click the         icon to add the objects to the Selected Objects section.         </li> <li>Note DTS cannot migrate data from the admin, local, or config database.</li> <li>A migration object can be a database, collection, or function.</li> <li>By default, the name of an object remains unchanged after migration. You can change the names of the objects in the destination instance by using the object name mapping feature. For more information, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time range for failed connections to the source or destination database	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time range based on your needs. If DTS reconnects to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>ONCE</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

9. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 10. After the data migration task passes the precheck, click Next.
- 11. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and read and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 12. Click **Buy and Start** to start the data migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

(?) Note We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the ApsaraDB for MongoDB instance.

- a. Wait until **Incremental Data Migration** and **The data migration task is not delayed** are displayed in the progress bar of the data migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the state of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

Task ID/Name:	Status: Migrating	Pause Task   View Details   Duplicate Task   Upgrade   Configure Monitoring and Alerting
2 2019-07-04 15:57:43 Created Full Data Migration 100 3 grated Rows: 5000003)		Completed Incremental Data Migration The migration task is not delayed.
Start Pause Stop Delete		Total: 1 item(s) , Per Page: 20 item(s) < < 1 > >

13. Switch your workloads to the destination ApsaraDB for MongoDB instance.

# 8.Migrate data between instances of the same Alibaba Cloud account

# 8.1. Migrate data between RDS instances

This topic describes how to use Data Transmission Service (DTS) to migrate data between RDS instances. DTS supports schema migration, full data migration, and incremental data migration. When you configure a data migration task, you can select all of the supported migration types to ensure service continuity.

#### Prerequisites

The database types of the RDS instances meet the following requirements.

Source database	Destination database
ApsaraDB RDS for MySQL ApsaraDB RDS for MariaDB TX	ApsaraDB RDS for MySQL ApsaraDB RDS for MariaDB TX
ApsaraDB RDS for SQL Server	ApsaraDB RDS for SQL Server
ApsaraDB RDS for PostgreSQL	ApsaraDB RDS for PostgreSQL

#### Precautions

- Data migration does not affect the data of the source database. During data migration, DTS reads the data of the source database and copies the data to the destination database. DTS does not delete the data of the source database. For more information, see Design concept of data migration.
- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- To ensure data consistency, we recommend that you do not write data to the source RDS instance during full data migration.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance will overwrite the data in the destination instance after the task is resumed.
- DTS automatically creates a database in the destination RDS instance. However, if the name of the

source database is invalid, you must manually create a database in the destination RDS instance before you configure the data migration task.

**Note** For more information about the naming conventions of ApsaraDB RDS and how to create a database, see Create databases and accounts for an ApsaraDB RDS for MySQL instance.

 If you migrate data between ApsaraDB RDS for PostgreSQL instances, take note of the following limits: After your workloads are switched to the destination database, newly written sequences do not increment from the maximum value of the sequences in the source database. Therefore, you must query the maximum value of the sequences in the source database before you switch your workloads to the destination database. Then, you must specify the queried maximum value as the starting value of the sequences in the destination database. You can execute the following statements to query the maximum value of the sequences in the source database:

```
do language plpgsql $$
declare
  nsp name;
  rel name;
  val int8;
begin
  for nsp,rel in select nspname,relname from pg_class t2 , pg_namespace t3 where t2.relna
mespace=t3.oid and t2.relkind='S'
  loop
     execute format($_$select last_value from %I.%I$_$, nsp, rel) into val;
     raise notice '%',
     format($_$select setval('%I.%I'::regclass, %s);$_$, nsp, rel, val+1);
   end loop;
end;
$$;
```

#### Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.	

#### **Migration types**

• Schema migration

DTS migrates the schemas of the required objects from the source RDS instance to the destination RDS instance.

• Full data migration

DTS migrates historical data of the required objects from the source RDS instance to the destination RDS instance.

Incremental data migration

After full data migration is completed, DTS synchronizes incremental data from the source RDS instance to the destination RDS instance. Incremental data migration allows you to ensure service continuity when you migrate data between RDS instances.

# SQL operations that can be synchronized during incremental data migration

Scenario	Operation type	SQL statement
• Migrate data	DML	INSERT, UPDATE, DELETE, and REPLACE
<ul> <li>between ApsaraDB RDS for MySQL instances</li> <li>Migrate data between ApsaraDB RDS for MariaDB TX instances</li> <li>Migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance</li> </ul>	DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREAT E FUNCT ION, CREAT E INDEX, CREAT E PROCEDURE, CREAT E TABLE, and CREAT E VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCAT E TABLE</li> </ul>
		INSERT, UPDATE, and DELETE
Migrate data between ApsaraDB RDS for SQL Server instances	DML	<b>Note</b> If an UPDATE operation updates only the large fields, DTS does not synchronize the operation.
	DDL	<ul> <li>ALTER TABLE, including only ADD COLUMN, DROP COLUMN, and RENAME COLUMN</li> <li>CREATE TABLE and CREATE INDEX</li> </ul>
		<b>Note</b> If a CREATE TABLE operation creates a partitioned table or a table that contains functions, DTS does not synchronize the operation.
		<ul> <li>DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>
Migrate data between	DML	INSERT, UPDATE, and DELETE
PostgreSQL instances Migrate data between ApsaraDB RDS for PPAS instances		

#### Permissions required for database accounts

#### Data Migration Migrate data betwe en instances of the same Alibaba Cl oud account

Scenario	Database	Schema migration	Full data migration	Incremental data migration
<ul> <li>Migrate data between ApsaraDB RDS for MySQL instances</li> </ul>	Source instance	The SELECT permission	The SELECT permission	The REPLICATION SLAVE, REPLICATION CLIENT, SHOW VIEW, and SELECT permissions
<ul> <li>Migrate data between ApsaraDB RDS for MariaDB TX instances</li> <li>Migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance</li> </ul>	Destinati on instance	The read and write permissions	The read and write permissions	The read and write permissions
	Source inst <i>a</i> nce	The SELECT permission	The SELECT permission	The owner permission on the objects to be migrated           Image: Note A privileged account has the required permissions.
	Destinati on instance	The read and write permissions	The read and write permissions	The read and write permissions
Migrate data between ApsaraDB RDS for SQL Server instances				

Scenario [	Database	Schema migration	Full data migration	Incremental data migration
Migrate data between ApsaraDB RDS for PostgreSQL instances	Source instance	The USAGE permission on pg_catalog	The SELECT permission on the objects to be migrated	rds_superuser Note A standard account of an ApsaraDB RDS for PostgreSQ L instance has the required permission s. If you receive a message indicating that the database account does not have the permission s of the superuser role, you must upgrade the kernel version of the RDS instance.

Scenario	Database	Schema migration	Full data migration	Incremental data migration	
		The CREATE and ti USAGE permissions on e the objects to be migrated	The permissions of the database owner, including the permissions to perform the INSERT, UPDATE, and DELETE operations	The permissions of the database owner, including the permissions to perform the INSERT, UPDATE, and DELETE operations	
	Destinati on instance		Note A standard account of an ApsaraDB RDS for PostgreSQL instance has the required permissions.	Note A standard account of an ApsaraDB RDS for PostgreSQL instance has the required permissions.	

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the **Migration Tasks** page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination	2.Configure Migration Types and Objects	3.Map name modification > 4.Precheck
* Task Name: R	DS_TO_RDS	
Source Database		
* Instance Type:	RDS Instance	
* Instance Region:	Singapore	,
* RDS Instance ID:	laisia Waing	RDS Instances of Other Apsara Stack Accounts
* Database Account:	dtstest	]
* Database Password:	••••••	→ Test Connectivity ⊘ Passed
Destination Database		
* Instance Type:	RDS Instance	
* Instance Region:	Singapore	
* RDS Instance ID:	-	
* Database Account:	dtstest	]
* Database Password:	•••••••	Test Connectivity
		Cancel Set Whitelist and Next

Section	Parameter	Description		
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to specify a unique task name.		
	Instance Type	Select RDS Instance.		
	Instance Region	Select the region where the source RDS instance resides.		
		Select the ID of the source RDS instance.		
	RDS Instance ID	<b>Note</b> The source and destination RDS instances can be the same or different. You can use DTS to migrate data within an RDS instance or between two RDS instances.		
		Enter the name of the source database in the ApsaraDB RDS for PostgreSQL instance.		
	Dat abase Name	<b>Note</b> This parameter is required only if the database engine of the RDS instance is <b>PostgreSQL</b> .		
	Dat abase Account	Enter the database account of the source RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.		
Source		Enter the password of the database account.		
Database	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		
		Select Non-encrypted or SSL-encrypted. If you want to select		
	Encryption	<ul> <li>SSL-encrypted, you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption for an ApsaraDB RDS for MySQL instance.</li> <li>Note         This parameter is required only if the database engine of the RDS instance is MySQL.         The Encryption parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.     </li> </ul>		

Section	Parameter	Description		
	Instance Type	Select RDS Instance.		
	Instance Region	Select the region where the destination RDS instance resides.		
		Select the ID of the destination RDS instance.		
	RDS Instance ID	<b>Note</b> The source and destination RDS instances can be the same or different. You can use DTS to migrate data within an RDS instance or between two RDS instances.		
	The name of	Enter the name of the destination database in the ApsaraDB RDS for PostgreSQL instance. The name of the destination database can be different from the name of the source database.		
	The name of the database.	<b>Note</b> This parameter is required only if the database engine of the RDS instance is <b>PostgreSQL</b> .		
Destinatio n Database	Dat abase Account	Enter the account that is used to connect to the RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
		Enter the password of the database account.		
	Dat abase Password	<b>?</b> Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		
		Select Non-encrypted or SSL-encrypted. If you want to select		
		<b>SSL-encrypted</b> , you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption for an ApsaraDB RDS for MySQL instance.		
	Encryption	Note This parameter is required only if the database engine of the RDS instance is MySQL. The Encryption parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.		

6. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

#### Data Migration Migrate data betwe en instances of the same Alibaba Cl oud account

1.Configu	re Source and   2.Configure Migration Types and  3.Advanced Settings   4.Precheck
* Migration Ty triggers, For mo	rpes: 🗹 Schema Migration 🛛 Full Data Migration 📝 Incremental Data Migration Note: Incremental data migration does not support
Note: do not cleans up th	clean up the incremental data log generated by the source database after the DTS task is started when the DTS full task is running. If the source database log too early, the DTS incremental task may fail
Data migrati between Aps For long-terr	on applies to short-term migration scenarios. Typical scenarios include migrating data to the cloud, scaling and sharding databases, and migrating data ara Stack databases. n data synchronization in real time, use the data synchronization feature.
Available	Selected (To edit an object name or its filter, hover over the object and click Edit.) Learn more.
Expand the	tree before you perform a glol   Q
🖃 💼 dtste	estdata 💿 dtstestdata (20bjects)
⊡ <u>⊫</u> 1a ⊕ <b>⊆</b> V	ews Extension Extensio Extension Extension Extension Extension Extension Extension Ext
	> er order
Select All	Remove All
* Reby Time for <b>Information:</b> 1. Data migration in the source do	r Failed Connection 720 Minutes () void on only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema itabase. b) operation during structure and full migration, otherwise the task may fail
	Cancel Previous Save Precheck
Setting	Description
	Select the migration types based on your business requirements. The migration types must be supported by the database engine.
	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> </ul>
Select the migratio	<ul> <li>To ensure service continuity during data migration, select Schema Migration, Full</li> <li>Data Migration, and Incremental Data Migration.</li> </ul>
n types	<b>Note</b> If <b>Incremental Data Migration</b> is not selected, we recommend that you do not write data to the source RDS instance during data migration. This ensures data consistency between the source and destination instances.

#### Dat a Transmission Service

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated. If you select tables or columns as the objects to be migrated, DTS does not migrate other objects such as views, triggers, and stored procedures to the destination database.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti ons to the source or destinati on	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time range based on your business requirements. If DTS reconnects to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
databas e	

#### 8. Click Precheck.

#### ? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the



icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 9. After the data migration task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

#### • Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

(?) Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of incremental data migration changes to The data migration task is not delayed again. Then, manually stop the migration task.



### 8.2. Migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you configure a data migration task, you can select all the supported migration types to ensure service continuity.

#### Prerequisites

An ApsaraDB RDS for MySQL instance is created. For more information, see Create an ApsaraDB RDS for MySQL instance.

**?** Note The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total size of the data in the ApsaraDB RDS for MariaDB TX instance.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source instance overwrites the data in the destination instance after the task is resumed.

#### Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration Free of charge.		Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.	

#### **Migration types**

• Schema migration

DTS migrates the schemas of required objects to the destination database. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function. DTS does not support schema migration for events.

#### ? Note

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- DTS does not migrate user information. To call a view, stored procedure, or function of the destination database, you must grant read and write permissions to INVOKER.

#### • Full data migration DTS migrates the historical data of required objects to the destination database.

#### ⑦ Note

- During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.
- To ensure successful data migration, we recommend that you do not perform DDL operations on the source database during full data migration.

#### • Incremental data migration

After full data migration is complete, DTS retrieves binary log files from the source database, and migrates incremental data to the destination database in real time.

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statement
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

#### Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
ApsaraDB RDS for MariaDB TX instance	The SELECT permission	The SELECT permission	The REPLICATION SLAVE, REPLICATION CLIENT, SHOW VIEW, and SELECT permissions
ApsaraDB RDS for MySQL instance	Read and write permissions	Read and write permissions	Read and write permissions

For more information about how to create a database account and grant permissions to the account, see the following topics:

- ApsaraDB RDS for MariaDB TX instances: Create an account on an ApsaraDB RDS for MariaDB TX instance.
- ApsaraDB RDS for MySQL instances: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account for an ApsaraDB RDS for MySQL instance.

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Dest	ination 2.Configure M	igration Types and Objects $>$	3.Advanced Settings	A.Precheck
* Task Nam	e: MariaDB TO MvSOL			
1				
Source Database				
* Instance T	ype: RDS Instance	Ŧ	DTS support type	
* Instance Re	gion: China (Hangzhou)	Ŧ		
* RDS Instance	e ID: rm-: qc	•	RDS Instances of Other Apsara Stack Accou	ints
* Database Acco	ount: dtstest			
* Database Passw	vord:	Ф	Test Connectivity 😔 Passed	
Destination Database				
* Instance T	ype: RDS Instance	•		
* Instance Re	gion: China (Hangzhou)	Ŧ		
* RDS Instance	e ID: m-i n7	-		
* Database Acco	ount: dtstest			
* Database Passw	vord: ••••••	٩	Test Connectivity 📀 Passed	
* Encryp	tion:  Non-encrypted  SSL-encry	pted		
				Cancel Set Whitelist and Next
Section	Parameter	Description		
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.		
	Instance Type	Type Select <b>RDS Instance</b> .		
	Instance	The region where	the source ApsaraDB F	RDS for MariaDB TX

instance resides.

The ID of the source ApsaraDB RDS for MariaDB TX instance.

Region

ID

**RDS** Instance

Section	Parameter	Description		
Source Dat <i>a</i> base	Dat abase Account	The database account of the source ApsaraDB RDS for MariaDB TX instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
	Dat abase Password	The password of the database account. After you specify the source database parameters, click <b>Test</b> <b>Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. <b>Output</b> If the specified parameters are valid, the <b>Passed</b> message is displayed. If the <b>Failed</b> message is displayed, click <b>Check</b> part to <b>Failed</b> Modify the source database		
		parameters based on the check results.		
	Instance Type	Select RDS Instance.		
	Instance Region	The region where the destination ApsaraDB RDS for MySQL instance resides.		
	RDS Instance ID	The ID of the destination ApsaraDB RDS for MySQL instance.		
	Dat abase Account	The database account of the destination ApsaraDB RDS for MySQL instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
		The password of the database account. After you specify the destination database parameters, click <b>Test</b> <b>Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid.		
Destination	Database Password	<b>Note</b> If the specified parameters are valid, the <b>Passed</b> message is displayed. If the <b>Failed</b> message is displayed, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.		
Instance Details				

Section	Parameter	Description		
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the ApsaraDB RDS for MySQL instance before you configure the data migration task. For more information, see <b>Configure SSL</b> encryption for an ApsaraDB RDS for MySQL instance.		
		<b>?</b> Note The Encryption parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.		

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### Q Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select migration types and the objects to migrate.

1.Configu	ure Source and	<ul> <li>2.Configure Migratio</li> </ul>	n Types and	3.Advanced Settings	$\rightarrow$	4.Precheck
* Migration 1 triggers. For m	iypes: 🗹 Schema Migra ore information, see Refe	tion 🔽 Full Data Migrat	ion 🔽 Incremental	Data Migration Note: Incre	mental data mig	ration does not support
Note: do no cleans up th	t clean up the increment ie log too early, the DTS	al data log generated by the s incremental task may fail	source database after th	e DTS task is started when the D	)TS full task is ru	nning. If the source database
Data migral between Ap For long-ter	ion applies to short-term sara Stack databases. m data synchronization ir	migration scenarios. Typical	scenarios include migrati hronization feature.	ing data to the cloud, scaling and	d sharding datab	ases, and migrating data
Available				Selected (To edit an object Edit.) Learn more.	name or its filter	r, hover over the object and dick
Expand the	e tree before you perform	a gloi 🔰 🔍				Q
🖃 📑 dtst	estdata			📼 dtstestdata (20)	piects)	
	ables (i.e			customer	Jecc)	
🗄 🥁 Views		> <	in order			
Select All				Remove All		
*Rename Data	abases and Tables:	Do Not Change Data	abase and Table Names	<ul> <li>Change Database and T</li> </ul>	able Names	
<ul> <li>Retry Time f</li> <li>Source table want to copy t the target data</li> <li>Information:</li> <li>Data migrat</li> <li>The source of 2. Do not do D</li> </ul>	DMS_ONLINE_Do you he temporary table to abase during DDL: ion only copies the data a latabase. DL operation during stru	Yes No (	abase and saves the cop rwise the task may fail	y in the destination database. Tl	he process does	not affect any data or schema
				Cancel	Previous	Save Precheck
etting	Description					
<ul> <li>To perform only full migration, select Schema Migration and Full Data Migration,</li> <li>To ensure service continuity during data migration, select Schema Migration,</li> <li>Select Data Migration, and Incremental Data Migration.</li> </ul>				Data Migration. Iigration, Full		
he nigratio types	Vou do no	e If Incremen t write data to tl	tal Data Mig he source data the source and	<b>ration</b> is not select abase during full da	ted, we re ata migrat	ecommend that ion. This ensures

#### Dat a Transmission Service

Setting	Description			
Select the objects to be migrate d	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.			
	<ul> <li>Notice</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>			
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.			
Specify the retry time for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.			
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.			

Setting	Description				
Specify whether to copy tempora ry tables to the destinati on dat abas e when DMS perform s online DDL operatio ns on the source table	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.				
	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.				
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.				
	<b>?</b> Note If you select No, the tables in the destination database may be locked.				

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

#### Stop the migration task

**Warning** We recommend that you prepare a rollback solution to migrate incremental data from the destination database to the source database in real time. This allows you to minimize the negative impact of switching your workloads to the destination database. For more information, see Switch workloads to the destination database. If you do not need to switch your workloads, you can perform the following steps to stop the migration task.

• Full data migration

Do not manually stop a task during full data migration. Otherwise, the system may fail to migrate all data. Wait until the migration task automatically ends.

• Incremental data migration

The task does not automatically end during incremental data migration. You must manually stop the migration task.

- i. Wait until the task progress bar shows **Incremental Data Migration** and **The migration task is not delayed**. Then, stop writing data to the source database for a few minutes. In some cases, the progress bar shows the delay time of **incremental data migration**.
- ii. After the status of incremental data migration changes to The migration task is not delayed, manually stop the migration task.



#### What's next

The database accounts that are used for data migration have read and write permissions. After data migration is complete, you must delete the database accounts of both the ApsaraDB RDS for MariaDB TX and ApsaraDB RDS for MySQL instances to ensure database security.

# 8.3. Migrate data from an ApsaraDB RDS for PostgreSQL instance to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from an ApsaraDB RDS for PostgreSQL instance to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration. When you configure a data migration task, you can select both of the supported migration types to ensure service continuity.

#### Prerequisites

- An ApsaraDB RDS for Post greSQL instance is created. For more information, see Create an ApsaraDB RDS for Post greSQL instance.
- An ApsaraDB RDS for MySQL instance is created. For more information, see Create an ApsaraDB RDS for MySQL instance.
- The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total size of the data in the ApsaraDB RDS for PostgreSQL instance.

#### Precautions

• DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when
the CPU utilization of the source and destination databases is less than 30%.

- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- To ensure that the data migration task runs as expected, you can perform a primary/secondary switchover only on a V11 ApsaraDB RDS for PostgreSQL instance. In this case, you must set the rds\_failover\_slot\_mode parameter to sync . For more information, see Logical Replication Slot Failover.

Q Warning If you perform a primary/secondary switchover on a self-managed PostgreSQL database or an ApsaraDB RDS for PostgreSQL instance of a version other than V11, the data migration task stops.

- If you modify the endpoint or zone of the ApsaraDB RDS for PostgreSQL instance, the data migration task stops.
- Before you configure a data migration task, you must create the corresponding databases and tables in the destination instance.
- A data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each database.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.
- If you perform a primary/secondary switchover during data migration, incremental data fails to be migrated. We recommend that you perform the primary/secondary switchover after the data migration task is complete.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba
Incremental data migration	Charged. For more information, see Pricing.	Pricing.

# **Migration** types

Migration type	Description
	DTS migrates the historical data of required objects from the source database to the destination ApsaraDB RDS for MySQL instance.
Full data migration	<b>Notice</b> During full data migration, we recommend that you do not perform DDL operations on the required objects. Otherwise, the objects may fail to be migrated.

Migration type	Description
Incremental data migration	DTS retrieves redo log files from the source database. Then, DTS migrates incremental data from the source database to the destination ApsaraDB RDS for MySQL instance. DTS can synchronize DML operations, such as INSERT, UPDATE, and DELETE. DTS cannot synchronize DDL operations. Incremental data migration allows you to ensure service continuity when you perform data migration.

# Permissions required for database accounts

Database	Full data migration	Incremental data migration
		Permissions of the superuser role
ApsaraDB RDS for PostgreSQL instance	The SELECT permission on the objects to migrate	<b>Note</b> If the source database runs on an ApsaraDB RDS for PostgreSQL V9.4 instance and you migrate only DML operations, the database account must have the REPLICATION permission.
ApsaraDB RDS for MySQL instance	Read and write permissions	Read and write permissions

For more information about how to create a database account and grant permissions to the account, see the following topics:

- ApsaraDB RDS for PostgreSQL instances: Create an account on an ApsaraDB RDS for PostgreSQL instance
- ApsaraDB RDS for MySQL instances: Create an account on an ApsaraDB RDS for MySQL instance
  - 1. Log on to the DTS console.
  - 2. In the left-side navigation pane, click **Data Migration**.
  - 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
  - 4. In the upper-right corner of the page, click **Experience the new console**.
  - 5. On the **Create Task** page, configure the data migration task.
    - i. Specify the information about the source and destination databases.

Section	Parameter	Description
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.
	Select Template	Select an existing database connection template for quick configuration.

Section	Parameter	Description	
	Database Type	Select PostgreSQL.	
	Access Method	Select Cloud Instance.	
Source	Instance Region	The region where the ApsaraDB RDS for PostgreSQL instance resides.	
Database	The ID of the instance.	The ID of the ApsaraDB RDS for PostgreSQL instance.	
	Database Name	The name of the ApsaraDB RDS for PostgreSQL instance.	
	Dat abase Account	The database account of the ApsaraDB RDS for PostgreSQL instance.	
	Database Password	The password of the database account.	
	Save as Template	Save the source database settings as a template.	
	Database Type	Select MySQL.	
Destinatio n Database	Access Method	Select Cloud Instance.	
	Instance Region	The region where the ApsaraDB RDS for MySQL instance resides.	
	RDS Instance ID	The ID of the ApsaraDB RDS for MySQL instance.	
	Dat abase Account	The database account of the ApsaraDB RDS for MySQL instance.	
	Database Password	The password of the database account.	
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the ApsaraDB RDS for MySQL instance before you configure the data migration task. For more information, see Configure SSL encryption for an ApsaraDB RDS for MySQL instance.	
		<b>?</b> Note The Encryption parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.	

Section	Parameter	Description
	Save as Template	Save the source database settings as a template.

- ii. Click **Test Connectivity and Proceed** to verify whether the parameters specified for the source and destination databases are valid.
  - If the specified parameters are valid, the **Passed** message is displayed.
  - If the Failed message is displayed, click Check next to Failed. Modify the parameters of the source or destination database based on the check results.
- iii. Select objects for the task and perform advanced settings.

Sectio	on	Parameter	Description
			To ensure data consistency, we recommend that you select both <b>Full Data Migration</b> and <b>Incremental Data Migration</b> . If you select only <b>Full Data Migration</b> , take note of the following items:
			<ul> <li>If the ApsaraDB RDS for PostgreSQL instance has data updates during full data migration, the data updates may not be migrated to the ApsaraDB RDS for MySQL instance.</li> </ul>
		Task Stages	<ul> <li>If the ApsaraDB RDS for MySQL instance has less than 2 CPU cores, we recommend that you disable the slow query logging feature.</li> <li>For more information, see Disable slow query log to improve migration performance.</li> </ul>
			If you select only <b>Incremental Data</b> Migration, take note of the following items:
			<ul> <li>DTS migrates only the incremental data that is generated in the ApsaraDB RDS for PostgreSQL instance after the task is started.</li> </ul>
			<ul> <li>Triggers cannot be synchronized during incremental data migration. For more information, see Configure a data synchronization task for a source database that contains a trigger.</li> </ul>

Section	Parameter	Description
	Select the processing mode of conflicting tables	<ul> <li>Precheck and Report Errors: checks whether the destination instance is empty. If the destination instance is empty, the precheck is passed. If the instance is not empty, an error is returned during the precheck and the data synchronization task cannot be started.</li> <li>Ignore Errors and Proceed: skips the check for empty destination instances.</li> <li>Warning If you select Ignore Errors and Proceed, the data records in the source database overwrite the data records that have the same keys in the destination database. Proceed with caution.</li> </ul>
Basic Settings		

Section	Parameter	Description
	Select the objects to be synchronized	<ul> <li>Select one or more objects from the Available section and click the  icon to add the objects to the Selected section.</li> <li>⑦ Note <ul> <li>You can select columns, tables, or databases as the objects to migrate.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the object shat are migrated to the destination database. For more information, see Object name mapping feature to rename the object name dependent on the object may fail to be migrated.</li> </ul> </li> </ul>
	Set alerts	<ul> <li>Specify whether to set alerts for the data migration task.</li> <li>If you do not want to set alerts, select No.</li> <li>If you want to set alerts, select Yes.</li> </ul>
Advanced		

Section	Parameter	Description
	Specify the retry time range for a failed connection to the source or destination instance.	Valid values: 10 to 1440. We recommend that you set the retry time range to more than 30 minutes. If multiple DTS instances have the same source or destination database, the lowest value takes effect. For example, the retry time is set to 30 minutes for Instance A and 60 minutes for Instance B, DTS retries failed connections at an interval of 30 minutes.

- iv. Click Next: Save Task Settings and Precheck in the lower part of the page.
- v. Run a **precheck** before you start the data migration task. DTS migrates data only after the task passes the precheck. If the task fails to pass the precheck, click **View Details** next to each failed item. After you troubleshoot the issues based on the causes, you can run a precheck again.
- vi. After the precheck is complete, click **Next: Purchase Instance** in the lower part of the page.
- vii. On the **Purchase Instance** page, specify the **Specification** parameter and read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- viii. Click Buy and Start to start the data migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

- a. Wait until Incremental Data Migration and The data migration task is not delayed are displayed in the progress bar of the data migration task. Then, stop writing data to the source database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the state of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.

? Note T	o delete a data migratio	on task, yo	u can select the task and click $\Box$ in the
lower part of	the page, and click the	Delete	to release the task.

6. Switch your workloads from the ApsaraDB RDS for PostgreSQL instance to the ApsaraDB RDS for MySQL instance. For more information, see Switch workloads to the destination database.

# 8.4. Migrate data from an ApsaraDB RDS for MySQL instance to a PolarDB for MySQL cluster

PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. PolarDB is compatible with the MySQL database engine and features high availability, ease of use, and reliability. This topic describes how to migrate data from an ApsaraDB RDS for MySQL instance to a PolarDB for MySQL cluster by using Data Transmission Service (DTS).

# Prerequisites

A PolarDB for MySQL cluster is created. For more information, see Create a PolarDB for MySQL cluster.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.
- If a data migration task fails, DTS automatically resumes the task. Therefore, before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination database after the task is resumed.

## Limits

• DTS supports the following types of objects for schema migration: table, view, trigger, stored procedure, and function.

Onte During schema migration, DTS changes the value of the SECURITY attribute from DE FINER to INVOKER for views, stored procedures, and functions.

• DTS does not migrate user information from the source database. After data migration is complete, if you want to call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

# **Migration types**

DTS supports schema migration, full data migration, and incremental data migration. For more information, see Terms.

**Note** When you migrate data between PolarDB for MySQL clusters, you can select all of the supported migration types to ensure service continuity.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

# Permissions required for database accounts

Database	Required permission
ApsaraDB RDS for MySQL	Read permissions on the objects to migrate
PolarDB for MySQL	Read and write permissions on the objects to migrate

**Note** For more information about how to create and authorize a database account, see Create an account on an ApsaraDB RDS for MySQL instance and Create a database account.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

#### Dat a Transmission Service

### Data Migration Migrate data betwe en instances of the same Alibaba Cl oud account

1.Configure Source and Destination	2.Configure Migration Types and Objects	3.Map name modification	4.Precheck
Tark Marrie La			
Task Name: R	DS_TO_PolarDB		
Source Database			
		7	
* Instance Type:	RDS Instance		
* Instance Region:	Singapore	]	
	Singapore		
* RDS Instance ID:	-	RDS Instances of Other Apsara Stack Accounts	
# Databaco Accounts		1	
Database Account.	utstest		
* Database Password:	······ Ø	Test Connectivity 📀 Passed	
Destination Database			
* Instance Type:	PolarDB		
* Instance Region:	Singapore		
* PolarDB Instance ID:			
		1	
* Database Account:	dtstest		
* Database Password:		Test Connectivity @ Passed	
			Cancel Cat Whitelist and Next
			Cancel Set Whitelist and Next

Section	Parameter	Description		
None	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.		
	Instance Type	Select RDS Instance.		
	Instance Region	The region where the source ApsaraDB RDS for MySQL instance resides.		
	Dat abase Account	The database account of the source ApsaraDB RDS for MySQL instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
		The password of the database account.		
Source Database	Dat abase Password	<b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message is displayed. If the <b>Failed</b> message is displayed, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.		

Section	Parameter	Description		
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the ApsaraDB RDS instance before you configure the data synchronization task. For more information, see Configure SSL encryption for an ApsaraDB RDS for MySQL instance.		
	Lifetypeion	<b>Note</b> The Encryption parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.		
	Instance Type	Select PolarDB.		
	Instance Region	The region where the destination PolarDB cluster resides.		
	PolarDB Instance ID	The ID of the destination PolarDB for MySQL cluster.		
Destination	Dat abase Account	The database account of the destination PolarDB cluster. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
Database		The password of the database account.		
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message is displayed. If the Failed message is displayed, click Check next to Failed. Modify the destination database parameters based on the check results.		

6. In the lower-right corner of the page, click Set Whitelist and Next.

### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

1.Configu	ire Source and	2.Configure Migration Ty	pes and	3.Advanced Setting	s >	4.Precheck	
<ul> <li>Migration T</li> <li>triggers. For m</li> </ul>	ypes: 🗹 Schema Migration ore information, see Referen	Full Data Migration	✓ Incremental [	Data Migration Note: I	incremental data migr	ation does not support	
Note: do no cleans up th	t clean up the incremental d e log too early, the DTS inc	ata log generated by the souro remental task may fail	e database after the	DTS task is started when t	the DTS full task is rur	nning. If the source data	abase
Data migrat between Ap For long-ter	ion applies to short-term mi sara Stack databases. m data synchronization in re	gration scenarios. Typical scena al time, use the data synchroni	arios include migratin ization feature,	g data to the cloud, scaling	g and sharding databa	ases, and migrating data	а
Available				Selected (To edit an ot Edit.) Learn more.	oject name or its filter	, hover over the object	and click
Expand the	tree before you perform a	gloi I <b>Q</b>				Q	
🗆 💼 dtst	estdata			💿 dtstestdata (	20bjects)		
	ables			customer	,		
•	16113		>	order			
			/				
			<				
Select All							
				Remove All			
*Rename Databases and Tables:      O Not Change Database and Table Names     O Change Database and Table Names							
* Retry Time f	or Failed Connection	720 Minutes	0				
*Source table want to copy to	DMS_ ONLINE_ Do you he temporary table to base during DDL :	🔾 Yes ( No					
the target data	base during DDL.						
Information: 1. Data migrat	on only copies the data and	schema in the source database	e and saves the copy	in the destination databas	e. The process does r	not affect any data or so	chema
in the source d 2. Do not do D	atabase. DL operation during structu	e and full migration, otherwise	the task may fail				
				Cancel	Previous	Save Pred	heck
Setting	Description						
	<ul> <li>To perform</li> </ul>	only full migratio	n, select <b>Sc</b> l	hema Migratio	n and Full E	Data Migratio	on.
		ervice continuity d	luring data r	nigration solor	t Schema M	igration Eul	I.
Select	Data Miora	ation, and Increm	ental Data	Migration.		igración, rul	L
he		and merem	Sector Batt				
nigratio		If Incremental	Data Migr	ation is not sel	ected we re	commend th:	ət
ntypes	you do not	write data to the	source data	base during full	data migrat	ion. This ensu	res
	data consis	tency between the	e source and	destination da	tabases.		

### Dat a Transmission Service

Setting	Description
Select the objects to be migrate d	Select one or more objects from the <b>Available</b> section and click the <b>&gt;</b> icon to move the objects to the <b>Selected</b> section.
	<ul> <li>Notice</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify	
whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
the source or destinati on databas	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
e	

Setting	Description
Specify whether to copy tempora ry tables to the destinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.
on databas e when DMS perform s online DDL operatio ns on the source table	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
	<b>Note</b> If you select No, the tables in the destination database may be locked.
destinati on databas e when DMS perform s online DDL operatio ns on the source table	<ul> <li><b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.</li> <li><b>No:</b> DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.</li> <li><b>Note</b> If you select No, the tables in the destination database may be locked.</li> </ul>

### 8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



12. Switch your workloads to the destination PolarDB cluster.

# 8.5. Migrate data from an ApsaraDB RDS for PPAS instance to a PolarDB for Oracle cluster

This topic describes how to migrate data from an ApsaraDB RDS for PPAS instance to a PolarDB for Oracle cluster by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you configure a data migration task, you can select all of the supported migration types to ensure service continuity.

## Prerequisites

- A PolarDB for Oracle cluster is created. For more information, see Create a cluster.
- The available storage space of the PolarDB for Oracle cluster is larger than the total size of the data in the ApsaraDB RDS for PPAS instance.
- To migrate incremental data from an ApsaraDB RDS for PPAS instance, you must submit a ticket to grant the permissions of the superuser role to the database account.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.

- A single data migration task can migrate data from only one database. To migrate data from multiple databases, you must create a data migration task for each database.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination instance after the task is resumed.
- •
- After your workloads are switched to the destination database, newly written sequences do not increment from the maximum value of the sequences in the source database. Therefore, you must query the maximum value of the sequences in the source database before you switch your workloads to the destination database. Then, you must specify the queried maximum value as the starting value of the sequences in the database. You can execute the following statements to query the maximum value of the sequences in the source database:

```
do language plpgsql $$
declare
  nsp name;
  rel name;
  val int8;
begin
  for nsp,rel in select nspname,relname from pg_class t2 , pg_namespace t3 where t2.relna
mespace=t3.oid and t2.relkind='S'
  loop
     execute format($_$select last_value from %I.%I$_$, nsp, rel) into val;
     raise notice '%',
     format($_$select setval('%I.%I'::regclass, %s);$_$, nsp, rel, val+1);
   end loop;
end;
$$;
```

- When you migrate data from an ApsaraDB RDS for PPAS instance to a PolarDB for Oracle cluster, we recommend that you take the following suggestions:
  - Make sure that the specifications of the PolarDB for Oracle cluster is greater than or equal to the specifications of the ApsaraDB RDS for PPAS instance. This prevents slow SQL queries or memory overflows due to insufficient CPU and memory after the migration task is completed. For more information about the recommended PolarDB for Oracle cluster specifications, see Specifications of an ApsaraDB RDS for PPAS instance and a PolarDB for Oracle cluster.
  - Select the compute nodes that meet your business requirements. For more information about the number of connections and IOPS, see Specifications of compute nodes.
  - Use a cluster endpoint to connect to your application. The cluster endpoint allows you to implement automatic read/write splitting. Read requests are forwarded to the read-only nodes This reduces the load on the PolarDB for Oracle cluster. For more information about how to obtain a cluster endpoint, see View or apply for an endpoint.

# Specifications of an ApsaraDB RDS for PPAS instance and a PolarDB for Oracle cluster

Make sure that the specifications of the PolarDB for Oracle cluster is greater than or equal to the specifications of the ApsaraDB RDS for PPAS instance. This prevents slow SQL queries or memory overflows due to insufficient CPU and memory after the migration task is completed. The following table lists the recommended specifications of a PolarDB for Oracle cluster.

Specifications of an ApsaraDB RDS for PPAS instance		Recommended specifications of a PolarDB for Oracle cluster		
Instance type	CPU and memory	Instance type	CPU and memory	
rds.ppas.t1.small	1 core, 1 GB	polar.o.x4.medium	2 cores, 8 GB	
ppas.x4.small.2	1 core, 4 GB	polar.o.x4.medium	2 cores, 8 GB	
ppas.x4.medium.2	2 cores, 8 GB	polar.o.x4.medium	2 cores, 8 GB	
ppas.x8.medium.2	2 cores, 16 GB	polar.o.x4.large	4 cores, 16 GB	
ppas.x4.large.2	4 cores, 16 GB	polar.o.x4.large	4 cores, 16 GB	
ppas.x8.large.2	4 cores, 32 GB	polar.o.x4.xlarge	8 cores, 32 GB	
ppas.x4.xlarge.2	8 cores, 32 GB	polar.o.x4.xlarge	8 cores, 32 GB	
ppas.x8.xlarge.2	8 cores, 64 GB	polar.o.x8.xlarge	8 cores, 64 GB	
ppas.x4.2xlarge.2	16 cores, 64 GB	polar.o.x8.2xlarge	16 cores, 128 GB	
ppas.x8.2xlarge.2	16 cores, 128 GB	polar.o.x8.2xlarge	16 cores, 128 GB	
ppas.x4.4xlarge.2	32 cores, 128 GB	polar.o.x8.4xlarge	32 cores, 256 GB	
ppas.x8.4xlarge.2	32 cores, 256 GB	polar.o.x8.4xlarge	32 cores, 256 GB	
rds.ppas.st.h43	60 cores, 470 GB	polar.o.x8.8xlarge	64 cores, 512 GB	

# Migration types

Migration type	Description
Schema migration	DTS migrates the schemas of the required objects from the source database to the destination PolarDB cluster. DTS supports schema migration for the following types of objects: table, view, synonym, trigger, stored procedure, stored function, package, and user-defined type.
	Notice In this scenario, DTS is incompatible with triggers. If an object contains triggers, data may become inconsistent between the source and destination databases.

Migration type	Description
Full data migration	DTS migrates the historical data of the required objects from the source database to the destination PolarDB cluster.
	<b>Notice</b> During schema migration and full data migration, do not perform data definition language (DDL) operations on the objects to be migrated. Otherwise, the objects may fail to be migrated.
Incremental data migration	DTS retrieves redo log files from the source database. Then, DTS synchronizes incremental data from the source database to the destination PolarDB cluster. DTS can synchronize data manipulation language (DML) operations, such as INSERT, UPDATE, and DELETE. DTS cannot synchronize DDL operations. Incremental data migration allows you to ensure service continuity when you migrate data from an ApsaraDB RDS for PPAS instance to a PolarDB cluster.

# Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.	

# Permissions required for database accounts

Log on to the source Oracle database, create an account for data collection, and grant permissions to the account.

Database	Schema migration	Full data migration	Incremental data migration
ApsaraDB RDS for PPAS The read permissions		The read permissions	The permissions of the superuser role
PolarDB for Oracle cluster	The permissions of the schema owner	The permissions of the schema owner	The permissions of the schema owner

For more information about how to create an account and grant permissions to the account, see the following topics:

- To migrate incremental data from an ApsaraDB RDS for PPAS instance, you must submit a ticket to grant the permissions of the superuser role to the database account.
- PolarDB for Oracle cluster: Create database accounts

# Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.

- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

Section	Parameter	Description		
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to specify a unique task name.		
	Instance Type	Select User-Created Database Connected Over Express Connect, VPN Gateway, or Smart Access Gateway. You cannot select ApsaraDB RDS for PPAS as the instance type.		
	Instance Region	Select the region where the ApsaraDB RDS for PPAS instance resides.		
	Peer VPC	Select the ID of the virtual private cloud (VPC) that is connected to the source database iin the ApsaraDB RDS for PPAS instance.		
	Database Type	Select PPAS.		
	Version	Select the database engine version of the ApsaraDB RDS for PPAS instance.		
	IP address	Enter the private IP address of the ApsaraDB RDS for PPAS instance.		
	Port Number	Enter the service port number of the ApsaraDB RDS for PPAS instance. The default port number is <b>3433</b> .		
	Database	Enter the name of the source database in the ApsaraDB RDS for PPAS		
	Name	listance.		
Source Dat abase	Name Dat abase Account	Enter the database account of the ApsaraDB RDS for PPAS instance. For information about the permissions that are required for the database account, see Permissions required for database accounts.		
Source Dat <i>a</i> base	Name Database Account	Enter the database account of the ApsaraDB RDS for PPAS instance. For information about the permissions that are required for the database account, see Permissions required for database accounts. Enter the password of the database account.		
Source Database	Dat abase Account Dat abase Password	Enter the database account of the ApsaraDB RDS for PPAS instance. For information about the permissions that are required for the database account, see Permissions required for database accounts. Enter the password of the database account. Onte After you specify the information about the self- managed Oracle database, you can click Test Connectivity next to Database Password to check whether the information is valid. If the information is valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Then, modify the information based on the check results.		
Source Dat abase	Dat abase Account Dat abase Password	Enter the database account of the ApsaraDB RDS for PPAS instance. For information about the permissions that are required for the database account, see Permissions required for database accounts. Enter the password of the database account. <b>Once</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		

Section	Parameter	Description		
	Instance Type	Select PolarDB.		
	Instance Region	Select the region where the destination PolarDB for Oracle cluster resides.		
	PolarDB Instance ID	Select the ID of the destination PolarDB for Oracle cluster.		
	Database Name	Enter the name of the destination database.		
Destinatio n Database	Dat abase Account	Enter the database account of the destination PolarDB for Oracle cluster. For information about the permissions that are required for the database account, see Permissions required for database accounts.		
		Enter the password of the database account.		
	Dat abase Password	<b>?</b> Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		

6. In the lower-right corner of the page, click Set Whitelist and Next.

### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types and the objects to be migrated.

	Incremental L	Jata Migration	
Data migration applies to short-term migration scenarios. Typical scen between Apsara Stack databases. For long-term data synchronization in real time, use the data synchron	arios include migratir ization feature.	g data to the cloud, scaling and sharding databases, and migrating o	data
Available		Selected (To edit an object name or its filter, hover over the object Edit.) Learn more.	ect and click
If you search globally, please expand the <b>Q</b>		Q	
e 🥌 mytest		MYTEST Source Database TABLE3	
	>		
	<		
Select All		Remove All	
Select All Name batch change:   No Yes		Remove All	

### Data Migration Migrate data betwe en instances of the same Alibaba Cl oud account

Setting	Description			
	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> </ul>			
Select the migratio n types	<ul> <li>Notice</li> <li>If Incremental Data Migration is not selected, we recommend that you do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.</li> <li>During schema migration and full data migration, do not perform DDL operations on the objects to be migrated. Otherwise, the objects may fail to be migrated.</li> </ul>			
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>			
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.			

Setting	Description			
Specify the retry time range for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.			
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business requirements. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.			

### 8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can troubleshoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.

sk Name 🗸 Search by migration task name.	Search Sort: Default Sorting V Status: All V	Tag
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade   Configure Monitoring and Alerting   Change password   Edit Tag
4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed Incremental Data M 1 n The migration task is not delayed.
3 Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) <pre> &lt; 1 &gt; &gt;</pre>

# 8.6. Migrate data between PolarDB for MySQL clusters

PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. PolarDB is compatible with the MySQL database engine and features high availability, ease of use, and reliability. This topic describes how to migrate data between PolarDB for MySQL clusters by using Data Transmission Service (DTS).

**Note** A PolarDB for MySQL cluster of an earlier version cannot be upgraded to 8.0. However, you can create a PolarDB for MySQL cluster of version 8.0, and then migrate data to this cluster. Before you migrate data between different versions of PolarDB for MySQL clusters, we recommend that you create a pay-as-you-go PolarDB cluster to test the compatibility. After testing, you can release the cluster.

# Prerequisites

- The source and destination PolarDB for MySQL clusters are created. For more information, see Create a PolarDB for MySQL cluster.
- The binary logging feature is enabled for the source PolarDB for MySQL cluster. For more information, see Enable binary logging.

# Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination database after the task is resumed.

# Migration types

DTS supports schema migration, full data migration, and incremental data migration. For more information, see Terms.

**Note** When you migrate data between PolarDB for MySQL clusters, you can select all of the supported migration types to ensure service continuity.

# Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

# Permissions required for database accounts

Database	Required permissions
Source PolarDB for MySQL cluster	The read permissions on the objects to be migrated
Destination PolarDB for MySQL cluster	The read and write permissions on the objects to be migrated

**Note** For more information about how to create and authorize a database account, see Create a database account.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.

- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

PolarDB

Instance ID

Database

Account

1.Configure Source and Destina	tion 2.Configure Mig	ration Types and Objects >	3.Advanced Setting	s >	4.Precheck
* Task Name:	PolarDB MySQL				
Source Database					
* Instance Typ			DTS support type		
* Instance Regio	D: China (Hangzhou)		• Or 5 support type		
* PolarDB Instance I	D: nc-				
* Database Account	at: dtetast				
* Database Pacewor	d:		Test Connectivity		
Database Passwol		Ψ	Test Connectivity	C Passeu	
Destination Database					
* Instance Typ	e: PolarDB	•			
* Instance Regio	n: China (Hangzhou)	•			
* PolarDB Instance I	D: pc-	<b>•</b>			
* Database Accour	nt: dtstest				
* Database Passwor	'd: ••••••	Ø	Test Connectivity	⊘ Passed	
		-	,		
					Cancel Set Whitelist and Next
Section	Parameter	Description			
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.			
	Instance Type	Select PolarDB.			
	Instance Region	Select the region wh	ere the source	e PolarDB cli	uster resides.

Select the ID of the source PolarDB cluster.

Enter the database account of the source PolarDB cluster.

Source Database

Section	Parameter	Description				
	Dat abase Password	Enter the password of the database account. Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the source database parameters based on the check results.				
	Instance Type	Select PolarDB.				
	Instance Region	Select the region where the destination PolarDB cluster resides.				
	PolarDB Instance ID	Select the ID of the destination PolarDB cluster.				
Destination	Dat abase Account	Enter the database account of the destination PolarDB cluster.				
Database		Enter the password of the database account.				
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.				

### 6. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note DTS adds the CIDR blocks of DTS servers to the whitelists of the source and destination PolarDB clusters. This ensures that DTS servers can connect to the source and destination PolarDB clusters.

### 7. Select the migration types and the objects to be migrated.

1.Configu	re Source and	2.Configure Migration Ty	pes and	3.Advanced Setti	ngs	4.Precheck
* Migration Ty triggers. For mo	ypes: 🗹 Schema Migration ore information, see Reference	✓ Full Data Migration ce	✓ Incremental D	ata Migration Note	e: Incremental data migra	tion does not support
Note: do not cleans up the	clean up the incremental da e log too early, the DTS incre	ta log generated by the souro emental task may fail	e database after the	DTS task is started whe	en the DTS full task is run	ning. If the source database
Data migrati between Aps For long-terr	on applies to short-term mig ara Stack databases. n data synchronization in rea	ration scenarios. Typical scena Il time, use the data synchroni	rios include migratin zation feature,	g data to the cloud, sca	ling and sharding databa	es, and migrating data
Available				Selected (To edit an Edit.) Learn more.	object name or its filter,	hover over the object and click
Expand the	tree before you perform a g	ol I Q				Q
🗆 💼 dtste	estdata			dtstestdata	a (20bjects)	
	ables			customer	(,	
	10W3		>	order		
			/			
			<			
Select All				Remove All		
*Rename Data	bases and Tables:	Do Not Change Database	and Table Names	O Change Databas	e and Table Names	
* Retry Time for Failed Connection 720 Minutes 🕐						
*Source table [ want to copy the the target data	DMS_ ONLINE_ Do you the temporary table to base during DDL:	🔿 Yes 💿 No 🕐				
Information	-					
1. Data migratio	on only copies the data and s	schema in the source database	and saves the copy	in the destination data	base. The process does n	ot affect any data or schema
2. Do not do Di	DL operation during structure	and full migration, otherwise	the task may fail			
				Can	cel Previous	Save Precheck
Setting	Description					
	• To perform	only full migratio	n, select <b>Scl</b>	nema Migrat	ion and Full D	ata Migration.
Select	<ul> <li>To ensure se</li> <li>Data Migra</li> </ul>	ervice continuity d tion, and Increm	uring data r ental Data	nigration, sele Migration.	ect Schema Mi	gration, Full
ine migratio	Notice	If Incremental	Data Migr	<b>ation</b> is not s	elected, we rea	commend that
n types	you do not v data consist	vrite data to the s ency between the	source datal	base during fu destination c	ull data migrati databases.	on. This ensures

### Dat a Transmission Service

Setting	Description
Select the objects to be migrate d	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
	<ul> <li>Notice</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to the source or destinati on databas e	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

Setting	Description			
Specify whether to copy tempora ry tables to the destinati on databas e when DMS perform s online DDL operatio ns on the source table	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.			
	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.			
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.			
	<b>Note</b> If you select No, the tables in the destination database may be locked.			

### 8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.

Task Name    Search by migration task name.	Search Sort: Default Sorting V Status: All V	N Tag
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplica	te Task Upgrade Configure Monitoring and Alerting Change password Edit Tag
4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Incremental Data M 1 n The migration task is not delayed.
Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) <

12. Switch your workloads to the destination PolarDB cluster.

### What's next

The database accounts that are used for data migration have the read and write permissions. After data migration is complete, you must delete the database accounts to ensure security.

# 8.7. Migrate data from a PolarDB for MySQL cluster to an ApsaraDB RDS for MySQL instance

ApsaraDB RDS is a stable, reliable, and scalable online database service. ApsaraDB RDS provides a complete database solution that includes disaster recovery, data backup, data recovery, and data migration. This topic describes how to migrate data from a PolarDB for MySQL cluster to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS).

### Prerequisites

- The binary logging feature is enabled for the PolarDB for MySQL cluster. For more information, see Enable binary logging.
- An ApsaraDB RDS for MySQL instance is created. For more information, see Create an ApsaraDB RDS for MySQL instance.

### Precautions

• DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.

- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.
- If a data migration task fails, DTS automatically resumes the task. Therefore, before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination database after the task is resumed.

## Limits

• DTS supports the following types of objects for schema migration: table, view, trigger, stored procedure, and function.

Onte During schema migration, DTS changes the value of the SECURITY attribute from DE FINER to INVOKER for views, stored procedures, and functions.

• DTS does not migrate user information from the source database. After data migration is complete, if you want to call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

## **Migration types**

DTS supports schema migration, full data migration, and incremental data migration. For more information, see Terms.

**Note** When you migrate data between PolarDB for MySQL clusters, you can select all of the supported migration types to ensure service continuity.

# Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data Charged. For more information, see Pricing.		information, see Pricing.	

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE

Operatio n type	SQL statements
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

# Permissions required for database accounts

Database	Required permissions
PolarDB for MySQL	The read permissions on the objects to be migrated
ApsaraDB RDS for MySQL	The read and write permissions on the objects to be migrated

**Note** For more information about how to create and authorize a database account, see Create a database account and Create an account on an ApsaraDB RDS for MySQL instance.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the destination instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

### Data Migration Migrate data betwe en instances of the same Alibaba Cl oud account

1.Configure Source and Destinatio	n > 2.Configure Migration Types and Objects >	3.Advanced Settir	ngs 🔰		4.Precheck
* Task Name: p	olarDB_To_RDS				
Source Database					
* Instance Type:	PolarDB	DTS support type			
* Instance Region:	China (Hangzhou)	•			
* PolarDB Instance ID:	pc-				
* Database Account:	dtstest	]			
* Database Password:	••••••	Test Connectivity	⊘ Passed		
Destination Database					
		7			
* Instance Type:	RDS Instance				
* Instance Region:	China (Hangzhou)				
* RDS Instance ID:	rm-				
* Database Account:	dtstest	]			
* Database Password:	4	Test Connectivity	⊘ Passed		
* Encryption:	Non-encrypted OSL-encrypted				
				Cancel	Set Whitelist and Next

Section	Parameter	Description			
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.			
	Instance Type	Select PolarDB.			
	Instance Region	Select the region where the source PolarDB cluster resides.			
	PolarDB Instance ID	Select the ID of the source PolarDB cluster.			
	Dat abase Account	Enter the database account of the source PolarDB cluster. For more information about the permissions that are required for the account, see Permissions required for database accounts.			
Source		Enter the password of the database account.			
Database	Dat abase Password	Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the source database parameters based on the check results.			

Section	Parameter	Description			
	Instance Type	Select RDS Instance.			
	Instance Region	Select the region where the destination RDS instance resides.			
	Dat abase Account	Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.			
		Enter the password of the database account.			
Destination Database	Database Password	<b>?</b> Note After you specify the destination database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.			
		Select Non-oncrupted or SSL-oncrupted if you want to select			
	Encryption	Select Non-encrypted of SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption on an ApsaraDB RDS for MySQL instance.			
		<b>Note</b> The <b>Encryption</b> parameter is available only for regions in mainland China and the China (Hong Kong) region.			

### 6. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note DTS adds the CIDR blocks of DTS servers to the whitelists of the source PolarDB cluster and the destination RDS instance. This ensures that DTS servers can connect to the source PolarDB cluster and the destination RDS instance.

### 7. Select the migration types and the objects to be migrated.
1.Configu	re Source and	2.Configure Migration T	ypes and	3.Advanced Settings	>	4.Precheck
* Migration Ty triggers. For mo	ypes: 🗹 Schema Migration ore information, see Refere	Full Data Migration	✓ Incremental I	Data Migration Note: Incr	remental data mig	ration does not support
Note: do not cleans up the Data migrati between Aps For long-terr	clean up the incremental d a log too early, the DTS inc on applies to short-term mi ara Stack databases. n data synchronization in re	ata log generated by the sou remental task may fail gration scenarios. Typical sce ral time, use the data synchro	rce database after the narios include migratir nization feature.	DTS task is started when the ng data to the cloud, scaling a	DTS full task is rund sharding datab	nning. If the source database
Available				Selected (To edit an objected (To edit.) Learn more,	ct name or its filte	r, hover over the object and click
Expand the	tree before you perform a	gloi I Q				Q
🗆 💼 dtste	estdata obles			🖮 dtstestdata (20	Objects)	
⊡ <b>—</b> 10	iews			customer		
			> <	■ order		
Select All				Remove All		
*Rename Data	bases and Tables:	Do Not Change Database	se and Table Names	Change Database and	Table Names	
* Retry Time fo	r Failed Connection	720 Minute	es 🕐			
*Source table I want to copy th the target data	DMS_ONLINE_Do you e temporary table to base during DDL:	🔾 Yes 💿 No 🥐				
Information: 1. Data migrati in the source da 2. Do not do Di	on only copies the data and atabase. DL operation during structu	schema in the source databa re and full migration, otherwis	se and saves the copy se the task may fail	r in the destination database.	The process does	not affect any data or schema
				Cancel	Previous	Save Precheck
etting	Description					
	• To perform	only full migrati	on, select <b>Sc</b>	hema Migration	and Full	Data Migration.
elect	• To ensure s Data Migra	ervice continuity ation, and Increr	during data i nental Data	migration, select 9 a Migration.	Schema M	ligration, Full
migratio n types ONOTE If Incremental Data do not write data to the source da consistency between the source ar			Data Migrat rce database urce and dest	t <b>ion</b> is not select e during data mig ination database	ed, we rec ration. This s.	ommend that you s ensures data

#### Dat a Transmission Service

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the <b>&gt;</b> icon to move the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
ons to the source or destinati on databas e	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

Setting	Description
Specify whether to copy tempora ry tables to the destinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.
on databas e when	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.
perform s online DDI	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
operatio ns on the source table	<b>Note</b> If you select No, the tables in the destination database may be locked.
DDL operatio ns on the source table	<b>Note</b> If you select No, the tables in the destination database may be locked.

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.

Task Name   Search by migration task name.	Search Sort: Default Sorting V Status: All V	bg
Task ID/Name:	Status: Migrating Quick Diagnostics   Pause Task   View Details   Duplicate Task	Upgrade Configure Monitoring and Alerting Change password Edit Tag
4 Nov 2021, 09:59:01 Created Schema Migration 100%	Full Data Migration 100%(Migrated Rows: 1)	Completed Incremental Data M 1 n The migration task is not delayed.
Pause Stop Delete Edit Tag		Total: 1 item(s), Per Page: 20 item(s) <pre>« &lt; 1 &gt; »</pre>

12. Switch your workloads to the destination RDS instance.

### 8.8. Migrate data between PolarDB for Oracle clusters

PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. It is compatible with MySQL, PostgreSQL, and Oracle database engines. PolarDB provides superior performance in storage and computing to meet diverse requirements of enterprises. This topic describes how to migrate data between PolarDB for Oracle clusters by using Data Transmission Service (DTS).

#### Prerequisites

- The source and destination PolarDB for Oracle clusters reside in the China (Shanghai) region. This is because data migration between PolarDB for Oracle clusters is available only in the China (Shanghai) region.
- The tables to migrate from the source PolarDB for Oracle cluster contain primary keys or UNIQUE NOT NULL indexes.
- The value of the wal\_level parameter is set to *logical* for the source PolarDB for Oracle cluster. This setting ensures that logical encoding is supported in write-ahead logging (WAL). For more information, see Configure cluster parameters.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- A data migration task can migrate data from only a single database. To migrate data from multiple

databases, you must create a data migration task for each database.

• During incremental data migration, if you select a schema to migrate, take note of the following requirements: If you create a table in the schema or run the RENAME command to rename a table in the schema, you must run the ALTER TABLE schema.table REPLICA IDENTITY FULL; command before you write data to the table.

**?** Note Replace the schema and table in the preceding sample command with the actual schema name and table name.

• To ensure that the latency of incremental data migration is accurate, DTS adds a heartbeat table named <a href="https://dts\_postgres\_heartbeat">dts\_postgres\_heartbeat</a> to the source database. The following figure shows the schema of the heartbeat table.



• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination database after the task is resumed.

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

#### **Migration types**

Migration type	Description
	DTS migrates the schemas of required objects from the source database to the destination PolarDB cluster. DTS supports schema migration for the following types of objects: table, view, synonym, trigger, stored procedure, function, package, and user-defined type.
Schema migration	O Notice DTS does not support triggers. If an object contains triggers, data may become inconsistent between the source and destination databases.

Migration type	Description		
	DTS migrates the historical data of required objects from the source database to the destination PolarDB cluster.		
Full data migration	<b>Notice</b> During schema migration and full data migration, do not perform DDL operations on the objects to migrate. Otherwise, the objects may fail to be migrated.		
Incremental data migration	DTS retrieves redo log files from the source database. Then, DTS synchronizes incremental data from the source database to the destination PolarDB cluster. DT can synchronize DML operations, such as INSERT, UPDATE, and DELETE. DTS cannot synchronize DDL operations. Incremental data migration allows you to ensure service continuity when you migrate data between PolarDB for Oracle clusters.		

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the **Migration Tasks** page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases.

1.Configure Source and Destination	2.Configure Migration Types and Objects >	3.A	dvanced Settings	>		4.Precheck
* Task Name: Pr	olarDB-O					
Source Database						
* Instance Type:	PolarDB	•	DTS support type			
* Instance Region:	China (Hangzhou)	•				
* PolarDB Instance ID:	prige-consistent	•				
* Database Name:	dtstestdata					
* Database Account:	dtstest					
* Database Password:	••••••	<b>\$</b> >	Test Connectivity			
Destination Database						
* Instance Type:	PolarDB	•				
* Instance Region:	China (Hangzhou)	•				
* PolarDB Instance ID:	terine to the All Starboltz	•				
* Database Name:	dtstestdata					
* Database Account:	dtstest					
* Database Password:	•••••	<b>\$</b> >	Test Connectivity			
					Cancel	Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.
	Instance Type	Select PolarDB.
	Instance Region	The region where the source PolarDB cluster resides.
	PolarDB Instance ID	The ID of the source PolarDB for Oracle cluster.
	Database Name	The name of the source database.
Source Database	Dat abase Account	A privileged account of the source PolarDB cluster. For more information, see Create database accounts.
		The password of the database account.
	Dat abase Password	<b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
	Instance Type	Select <b>PolarDB</b> .
	Instance Region	The region where the destination PolarDB cluster resides.
	PolarDB Instance ID	The ID of the destination PolarDB for Oracle cluster.
	Dat abase Name	The name of the destination database.
		The database account of the destination PolarDB for Oracle cluster. The account must have permissions of the <b>database owner</b> .
Destinatio	Dat abase Account	<b>Notice</b> You can specify the <b>database owner</b> when you create a database.
n Database		
Bacabase		

Section Parar	meter	Description
		The password of the database account.
Dat al Passv	base word	<b>?</b> Note After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration type and the objects to be migrated.

Configure Source and Destination	2.Configure Migration T	ypes and	3.Advanced Settings	$\rightarrow$	4.Precheck
<ul> <li>Migration Types: Schema triggers. For more information, see</li> <li>Note: do not clean up the incre cleans up the log too early, the</li> </ul>	Migration Full Data Migration e Reference mental data log generated by the sou DTS incremental task may fail	☑ Incremental I rce database after the	Data Migration Note: Increm DTS task is started when the D	nental data migrati TS full task is runni	ion does not support ing. If the source database
Availabla			Selected (To add an object	nama or ita filtar h	naver over the object and dick
Expand the tree before you pe	rform a gloi 🔰 🔍		Edit.) Learn more.		
		> <			
Select All					
*Rename Databases and Tables: * Retry Time for Failed Connection	Do Not Change Databa n 720 Minut	ise and Table Names	Remove All Change Database and Tr	able Names	
*Rename Databases and Tables: * Retry Time for Failed Connection Information: 1. Data migration only copies the i in the source database. 2. Do not do DDL operation during	Do Not Change Databa n     720     Minul data and schema in the source datab g structure and full migration, otherw	ese and Table Names es ⑦ ase and saves the copy se the task may fail	Remove All  Change Database and Ta in the destination database. Th Cancel	able Names ne process does not Previous	t affect any data or schema
*Rename Databases and Tables: * Retry Time for Failed Connection Information: 1. Data migration only copies the u in the source database. 2. Do not do DDL operation during tting	Do Not Change Databa T20     Minul data and schema in the source datab g structure and full migration, otherw Description	ese and Table Names es ⑦ ase and saves the copy se the task may fail	Remove All  Cancel  Cancel	able Names e process does not Previous	t affect any data or schema
*Rename Databases and Tables: * Retry Time for Failed Connection Information: 1. Data migration only copies the i in the source database. 2. Do not do DDL operation during tting	Do Not Change Databa Tzo     Minul data and schema in the source datab g structure and full migration, otherw Description     To perform only Dat a Migratio	se and Table Names es ⑦ ase and saves the copy se the task may fail y full dat a min.	Remove All  Change Database and Tr  change Database and Tr  cancel  igration, select Sc	able Names e process does not Previous hema Mig	t affect any data or schema Save Precheck gration and Full
*Rename Databases and Tables: * Retry Time for Failed Connection Information: 1. Data migration only copies the v in the source database. 2. Do not do DDL operation during tting	Do Not Change Database     Tzo     Minul     data and schema in the source database     structure and full migration, otherw      Description     O To perform only     Data Migration     To ensure service     Migration, Full	es and Table Names es (2) ase and saves the copy se the task may fail y full data min. te continuity Data Migra	Remove All  Change Database and Tr  change Database and Tr  in the destination database. Th  cancel  igration, select Sc  during data migra  ition, and Increm	able Names e process does not Previous hema Mig tion, selec ental Dat	save Precheck Save Precheck gration and Full t Schema ta Migration.
*Rename Databases and Tables: * Retry Time for Failed Connection <b>Information:</b> 1. Data migration only copies the rinthe source database. 2. Do not do DDL operation during tting	<ul> <li>Do Not Change Database</li> <li>T20</li> <li>Minute</li> </ul>	es and Table Names es (2) ase and saves the copy se the task may fail y full data min n. te continuity Data Migra	Remove All  Change Database and Tr  the destination database. Th  cancel  igration, select Sc  during data migra  it ion, and Increm	able Names e process does not Previous hema Mig tion, selec ental Dat	save Precheck Save Precheck gration and Full t Schema ta Migration.
*Rename Databases and Tables: * Retry Time for Failed Connection Information: 1. Data migration only copies the inthe source database. 2. Do not do DDL operation during ttring	<ul> <li>Do Not Change Database</li> <li>To perform only Data Migration, otherw</li> <li>To perform only Data Migration, Full</li> <li>To ensure service Migration, Full</li> <li>Notice <ul> <li>If Increar that you data mide and des</li> </ul> </li> </ul>	ese and Table Names es (?) ase and saves the copy se the task may fail y full data min. the continuity Data Migra nental Data gration. This tination data	Remove All  Change Database and To  Cancel  igration, select Sc  during data migra  ition, and Increm  Migration is not e data to the sour ensures data cons bases.	able Names e process does not Previous hema Mig tion, selec ental Dat : selected, ce databa: istency be	se during full street environment se during full

Setting	Description		
	Select one or more objects from the <b>Available</b> section and click the > icon to add the objects to the <b>Selected</b> section. You can select columns, tables, or schemas as the objects to migrate.		
Select the objects that you want to migrate	<ul> <li>After an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination PolarDB cluster. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>		
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.		
Specify the retry time	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time range based on your needs. If DTS reconnects to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.		
connections to the source or destination database	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.		

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.

- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



#### References

Switch workloads to the destination database

# 8.9. Migrate data from a standalone instance to a replica set or sharded cluster instance

ApsaraDB for MongoDB provides standalone instances, replica set instances, and sharded cluster instances. Standalone instances are suitable in development environments and test environments that store unimportant data. Replica set instances and sharded cluster instances are suitable in production environments. This topic describes how to migrate data from a standalone instance to a replica set instance or sharded cluster instance by using Data Transmission Service (DTS).

#### Prerequisites

The available storage space of the destination instance is larger than the total size of the data in the source instance.

#### Precautions

• DTS uses specific resources of the source instance and destination instance during the full data

migration process. This may increase the loads of the physical hosts on which these instances are deployed. If the data volume is large or the specifications of the physical hosts do not meet your requirements, the source instance and destination instance may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.

- DTS does not support incremental data migration from a standalone instance. To ensure data consistency, do not write data to the source instance during the full data migration process.
- If the source and destination ApsaraDB for MongoDB instances have different versions or storage engines, make sure that the versions or storage engines are compatible. For more information, see MongoDB versions and storage engines.
- •
- DTS cannot migrate data from the admin or local database.

#### Pricing

Migration type	Instance fee	Internet traffic fee
Full data migration	Free of charge	Free of charge

#### Migration types

All existing data in the source instance is migrated to the destination instance. This is called full data migration.

Onte Full data migration is supported for databases, collections, and indexes.

#### Permissions required for database accounts

Instance	Permission
Source ApsaraDB for MongoDB instance	Read permissions on the source database
Destination ApsaraDB for MongoDB instance	Read and write permissions on the destination database

For more information about how to create and authorize a database account, see Manage MongoDB databases by using DMS.

#### Procedure

- 1.
- 2.
- 3. In the left-side navigation pane, click **Replica Set Instances**.
- 4.
- 5. In the upper-right corner of the page that appears, click Migrate Databases.
- 6. On the Create Migration Task page, create a migration task.
  - i. Specify the information about the source database and destination database.

1.Configure Source and Destination	2.Configure Migration Types and Objects 🔪	3.Map name modification > 4.Precheck
* Task Name:		
Source Database		
* Instance Type:	ApsaraDB for MongoDB	
* Instance Region:	China (Hangzhou)	
* MongoDB Instance ID:	dds-bp	
* Database Name:	admin	Authenticate Database with Account
* Database Account:	1981	
* Database Password:	•••••	Test Connectivity
Destination Database		
* Instance Type:	MongoDB Instance	
* Instance Region:	China (Hangzhou)	
* MongoDB Instance ID:	dds-bt	
* Database Name:	admin	Authenticate Database with Account
* Database Account:		
* Database Password:	•••••	Test Connectivity

Cancel Assess Data Migration to Cloud Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	The name of the migration task. DTS automatically generates a task name. We recommend that you specify an informative name that helps identify the migration task. The task name does not need to be unique.
	Instance Type	Select ApsaraDB for MongoDB.
	Instance Region	Select the region where the source ApsaraDB for MongoDB instance is deployed.
	MongoDB Instance ID	Select the ID of the source instance.
Dat abase Name		Enter the name of the destination database. The permissions on the destination database must be granted to the account that is used to connect to the destination instance.
		<b>Note</b> If you want to use the root account, enter admin in the Database Name field.
Source Database	Enter the username of the account that is used to connect to the source instance. For more information about the permissions the are required for the account, see Permissions required for datab accounts.	

Section	Parameter	Description
		Enter the password of the account that is used to connect to the source instance.
	Dat abase Password	<b>?</b> Note After you specify the information about the source database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is correct. If the information is correct, the <b>Passed</b> message appears. If the information is incorrect, the <b>Failed</b> message appears and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.
	Instance Type	Select MongoDB Instance.
	Instance Region	Select the region where the destination ApsaraDB for MongoDB instance is deployed.
	MongoDB Instance ID	Select the ID of the destination instance.
	Dat abase Name	Enter the name of the destination database. The permissions on the destination database must be granted to the account that is used to connect to the destination instance.
Destienti		<b>Note</b> If you want to use the root account, enter admin in the Database Name field.
on		Enter the username of the account that is used to connect to the
Database	Dat abase Account	destination instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the account that is used to connect to the destination instance.
	Database Password	<b>?</b> Note After you specify the information about the destination database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is correct. If the information is correct, the <b>Passed</b> message appears. If the information is incorrect, the <b>Failed</b> message appears and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.

#### ii. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note DTS automatically creates IP address whitelists on the source instance and destination instance. The created IP address whitelists contain the IP address of the DTS server. This ensures that the DTS server can communicate with the source instance and destination instance. After the data of the source instance is migrated to the destination instance, you can delete these IP address whitelists. For more information, see **Configure a** whitelist for a sharded cluster instance.

iii. Configure migration types and objects to be migrated.

1.Configure Source and	Destination 2.Configure Migration Types and O	ojects	3.Map name modification	$\rightarrow$	4.Precheck
* M	gration Types: I Full Data Migration				
Ava If	lable you search globally, please expand (	>	Selected (To edit an object name of and click Edit.) Learn more.	r its filter, hover over th	e object Hover over the required object and dick Edit. In the dalog box that appears, modify the object name of the destination database and select the columns to migrate.
Sele	ct All		Remove All		
* Nan chang 1. Da data 2. DC	e batch   No Yes e : mation: ta migration only copies the data and schema in the source data schema in the source database. L operations are not supported during data migration because	tabase and saves this can cause m	s the copy in the destination database, sigration failures.	The process does not af	fect any
			С	ancel Previous	Save Precheck
Paramet er	Description				
	Select Full Data Migration.				

Migration Types	Note If the source instance is a standalone instance, you can select only Full Data Migration. To ensure data consistency, do not write data to the source instance during the full data migration process.

Paramet er	Description
Objects	<ul> <li>Select objects to be migrated in the following way:</li> <li>a. Select one or more objects from the Available section.</li> <li>b. Click &gt; to move the selected objects to the Selected section.</li> <li>? Note</li> <li>DTS cannot migrate data from the admin or local database.</li> <li>The objects that you can migrate are databases, collections, and functions.</li> <li>By default, after an object is migrated to the destination instance, the</li> </ul>
Granifu	name of the object remains unchanged. If you want an object to have a different name after the migration, you can use the object name mapping feature of DTS to rename the object. For more information, see <mark>Object</mark> name mapping.
Specify whether to rename objects	You can use the object name mapping feature of DTS to rename the objects that you want to migrate. For more information, see Object name mapping.
Specify the retry time for a failed connecti	By default, if DTS fails to connect to the source instance or destination instance, DTS retries to establish a connection within the next 12 hours. You can specify the retry time based on your business requirements. If DTS reconnects to the source instance and destination instance within the specified period of time, DTS resumes the migration task. If DTS cannot reconnect to the source instance or destination instance within the specified period task fails.
source instance or destinati on instance	<b>Note</b> When DTS retries to establish a connection, you are charged for the migration task. We recommend that you specify the retry time based on your business requirements. After the source instance and destination instance are released, we recommend that you release the migration task at your earliest opportunity.

iv. In the lower-right corner of the page, click **Precheck**.

? Note

- Before you can start the data migration task, DTS performs a precheck. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed

item to view details.

- After you troubleshoot the issues based on the causes, you can run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- v. After the data migration task passes the precheck, click Next.
- vi. In the **Confirm Settings** dialog box, configure the **Instance Class** parameter. Then, read and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- vii. Click Buy and Start to start the data migration task.
- 7. View the progress of the migration task.

When the status of the migration task changes to **Completed**, the data of the source instance is migrated to the destination instance.

(?) Note Do not manually stop a migration task. If you manually stop a migration task, the data that is migrated to the destination instance is incomplete. You can wait until the data migration task automatically stops.

	Task ID/Name:	Status: Completed View Details	Duplicate Task	Configure Monitoring and Alerting Modify password Dec 16, 2019, 14:13:49 Completed
	Full Data Migration 100%(Migrated Rows: 677311)			

8. Switch your workloads to the destination ApsaraDB for MongoDB instance.

#### What to do next

If you no longer need the source instance, release the source instance.

- If the source instance is charged based on the pay-as-you-go billing method, you can release the instance. For more information, see Release an instance or a node.
- If the source instance is charged based on the subscription billing method, you are not allowed to release the instance.

#### References

- Connect to a replica set instance
- Connect to a sharded cluster instance
- Configure sharding to maximize the performance of shards

## 8.10. Migrate data from an ApsaraDB for MongoDB replica set instance to an ApsaraDB for MongoDB sharded cluster instance

This topic describes how to migrate data from an ApsaraDB for MongoDB replica set instance to an ApsaraDB for MongoDB sharded cluster instance by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration. When you migrate data between ApsaraDB for MongoDB instances, you can select both of the supported migration types to ensure service continuity.

#### Prerequisites

Each shard in the destination sharded cluster instance has sufficient storage space.

#### Precautions

- DTS consumes the resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If you migrate a large amount of data or if the server specifications do not meet your requirements, database services may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- If the source and destination ApsaraDB for MongoDB instances have different versions or storage engines, make sure that the versions or storage engines are compatible. For more information, see MongoDB versions and storage engines.

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba
Incremental data migration	Charged. For more information, see Pricing.	Cloud over the Internet. For more information, see Pricing.

#### **Migration** types

Migration type	Description
	DTS migrates the historical data of required objects from the source MongoDB database to the destination MongoDB database.
Full data migration	<b>Note</b> The following types of objects are supported: database, collection, and index.

Migration type	Description	
	After full data migration is complete, DTS synchronizes incremental data from the source MongoDB database to the destination MongoDB database.	
Incremental data migration	<ul> <li>Note</li> <li>The create and delete operations that are performed on databases, collections, and indexes can be synchronized.</li> <li>The create, delete, and update operations that are performed on documents can be synchronized.</li> </ul>	

#### Permissions required for database accounts

Instance	Full data migration	Incremental data migration
ApsaraDB for MongoDB replica set instance	Read permissions on the source database	Read permissions on the source database, the admin database, and the local database
ApsaraDB for MongoDB sharded cluster instance	Read and write permissions on the destination database	Read and write permissions on the destination database

Note For more information about how to create and authorize a database account, see Manage user permissions on MongoDB databases.

#### Preparations

Create databases and collections to be sharded in the destination ApsaraDB for MongoDB instance, and configure data sharding based on your business requirements. For more information, see Configure sharding to maximize the performance of shards.

**?** Note After you configure sharding for a cluster, the migrated data is distributed among different shards. This maximizes the performance of the sharded cluster.

#### Procedure

- 1.
- 2. In the left-side navigation pane, click **Replica Set Instances**.

3.

4.

- 5. In the upper-right corner of the page that appears, click **Migrate Databases**.
- 6. On the Create Migration Task page, create a migration task.
  - i. Configure the source and destination databases.

#### Dat a Transmission Service

#### Data Migration Migrate data betwe en instances of the same Alibaba Cl oud account

1.Configure Source and Destination	2.Configure Migration Types and Objects 🔪	3.Map name modification A.Precheck
* Task Name: (		
Source Database		
* Instance Type:	ApsaraDB for MongoDB	
* Instance Region:	China (Hangzhou)	
* MongoDB Instance ID:	dds-bp	
* Database Name:	admin	Authenticate Database with Account
* Database Account:		]
* Database Password:	•••••• Ø	Test Connectivity
Destination Database		
* Instance Type:	MongoDB Instance	
* Instance Region:	China (Hangzhou)	
* MongoDB Instance ID:	dds-bt	
* Database Name:	admin	Authenticate Database with Account
* Database Account:		]
* Database Password:	••••••	Test Connectivity

Cancel Assess Data Migration to Cloud Set Whitelist and Next

Section	Parameter	Parameter
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.
	Instance Type	Select ApsaraDB for MongoDB.
	Instance Region	The region where the source ApsaraDB for MongoDB instance resides.
	MongoDB Instance ID	The ID of the source instance.
		The name of the authentication database. The database account is created in this database.
Dat abase Name		<b>Note</b> If you want to use the root account, enter admin in the Database Name field.
Source	Dat abase Account	The username of the account that is used to connect to the source instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
Database		

Section	Parameter	Parameter		
		The password of the database account.		
	Dat abase Password	<b>?</b> Note After you specify the information about the source database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is correct. If the information is correct, the <b>Passed</b> message is displayed. If the information is incorrect, the <b>Failed</b> message is displayed and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.		
	Instance Type	Select MongoDB Instance.		
	Instance The region where the destination ApsaraDB for MongoDB in Region resides.			
	MongoDB Instance ID	The ID of the destination instance.		
	Dat abase Name	The name of the authentication database. The database account is created in this database.		
		<b>Note</b> If you want to use the root account, enter admin in the Database Name field.		
Destinati on				
on Database	Dat abase Account	The username of the account that is used to connect to the destination instance. For information about the permissions that are required for the account, see Permissions required for database accounts.		
		The password of the database account.		
	Database Password	<b>?</b> Note After you specify the information about the destination database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is correct. If the information is correct, the <b>Passed</b> message appears. If the information is incorrect, the <b>Failed</b> message appears and you must click <b>Check</b> next to the <b>Failed</b> message to modify the information.		

#### ii. In the lower-right corner of the page, click Set Whitelist and Next.

#### û Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.

	> <		database and select columns to migrate.
Select All			

#### iii. Configure migration types and the objects to migrate.

Setting	Description
Select migrati on types	<ul> <li>If you want to migrate only the existing data in the source instance, select Full Data Migration.</li> <li>If you want to migrate data without downtime, select Full Data Migration and Incremental Data Migration.</li> <li>Note If Incremental Data Migration is not selected, do not write data to the source instance during full data migration. This ensures data consistency between the source instance and destination instance.</li> </ul>
	<ul> <li>The method to select objects to migrate:</li> <li>a. Select one or more objects from the Available section.</li> <li>b. Click &gt; to add the selected objects to the Selected section.</li> </ul>
Select objects to migrate	<ul> <li>Note</li> <li>DTS cannot migrate data from the admin or local database.</li> <li>The objects that you can migrate are databases, collections, and functions.</li> <li>By default, after an object is migrated to the destination instance, the name of the object remains unchanged. If you want to rename an object in the destination database, you can use the object name mapping feature. For more information, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time range for a failed connect ion to the source or destinat	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails. <b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
ion databas e	

iv. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, DTS performs a precheck. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed

item to view details.

- After you troubleshoot the issues based on the causes, you can run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- v. After the data migration task passes the precheck, click Next.
- vi. In the **Confirm Settings** dialog box, configure the **Instance Class** parameter. Then, read and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- vii. Click **Buy and Start** to start the data migration task.
- 7. View the progress of the migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

- Incremental data migration An incremental data migration task does not automatically stop. You must manually stop the task.
  - i. Wait until **Incremental Data Migration** and **The data migration task is not delayed** appear in the progress bar of the data migration task. Then, stop writing data to the source instance for a few minutes. The latency of **incremental data migration** may be displayed in the progress bar.
  - ii. Wait until the state of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.
- 8. Switch your workloads to the destination ApsaraDB for MongoDB instance.

#### What to do next

If you no longer need the source instance, release the source instance.

- If the source instance is charged based on the pay-as-you-go billing method, you can release the instance. For more information, see Release an instance or a node.
- If the source instance is charged based on the subscription billing method, you are not allowed to release the instance.

## 8.11. Migrate data between ApsaraDB for MongoDB instances across regions

This topic describes how to migrate a standalone instance or a replica set instance across regions by using Data Transmission Service (DTS). You can perform full data migration and incremental data migration for a replica set instance. You can perform only full data migration for a standalone instance.

#### Prerequisites

• The source instance is a standalone instance or a replica set instance. If the source instance is a sharded cluster instance, we recommend that you use the built-in tools of MongoDB to migrate data. For more information, see Migrate a self-managed MongoDB database to ApsaraDB for MongoDB by using tools provided by MongoDB.

**?** Note DTS does not support incremental data migration from a standalone instance. For more information, see Migration types.

- A public endpoint is applied for the source instance. For more information, see (Optional) Apply for a public endpoint for an ApsaraDB for MongoDB instance.
- The destination instance is created in the destination region. For more information, see Create a standalone instance, Create a replica set instance, or Create a sharded cluster instance.

(?) Note The available storage space of the destination instance must be larger than the total size of the data in the source instance.

#### Context

You may need to migrate the data of an ApsaraDB for MongoDB instance across regions in the following situations:

- Restructure your business.
- Use an ApsaraDB for MongoDB instance to provide database services for applications deployed on an Elastic Compute Service (ECS) instance. However, the ApsaraDB for MongoDB instance and the ECS instance are not in the same region.

In this example, an ApsaraDB for MongoDB instance in the China (Qingdao) region is migrated to an ApsaraDB for MongoDB instance in the China (Hangzhou) region.



**Note** This example describes only how to migrate data from the source instance. After data is migrated, you can release the source instance if you no longer need the source instance.

#### Precautions

• DTS consumes the resources of the source and destination instances during full data migration. This may increase the loads of the database servers. If you migrate a large volume of data or the server

specifications cannot meet your requirements, database services may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.

- DTS does not support incremental data migration from a standalone instance. To ensure data consistency, do not write data to the source instance during full data migration.
- If the source and destination MongoDB databases use different versions or storage engines, make sure that your applications can run on both databases. For more information about the versions and storage engines that are supported by ApsaraDB for MongoDB, see MongoDB versions and storage engines.

•

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba
Incremental data migration	Charged. For more information, see Pricing.	Cloud over the Internet. For more information, see Pricing.

#### **Migration types**

Migration type	Description
Full data migration	DTS migrates the historical data of required objects from the source MongoDB database to the destination MongoDB database.
	Note The following types of objects are supported: database, collection, and index.
	After full data migration is complete. DTS synchronizes incremental data from the
Incremental data migration	source MongoDB database to the destination MongoDB database.
	<ul> <li>Note</li> <li>The create and delete operations that are performed on databases, collections, and indexes can be synchronized.</li> <li>The create, delete, and update operations that are performed on documents can be synchronized.</li> </ul>

#### Permissions required for database accounts

Database	Full data migration	Incremental data migration
Source ApsaraDB for MongoDB instance	The read permissions on the source database	The read permissions on the source database, the admin database, and the local database
Destination ApsaraDB for MongoDB instance	The read and write permissions on the destination database	The read and write permissions on the destination database

Note For more information about how to create and authorize a database account, see Manage user permissions on MongoDB databases.

#### Procedure

- 1.
- 2.
- 3. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
- 4.
- 5. In the upper-right corner of the page that appears, click Migrate Databases.
- 6. On the Create Migration Task page, create a data migration task.
  - i. Configure the source and destination databases.

1.Configure Source a	nd Destinatior	n 💙 2.Configu	re Migration Types and Objects	>	3.Map name modification	>	4.Precheck
* T	ask Name: 🛛						
Source Database							
* Ir	nstance Type:	ApsaraDB for Mong	DB	Ŧ			
* Inst	tance Region:	China (Qingdao)		•			
* MongoDB	Instance ID:	dds-bp		•			
* Dat	tabase Name:	admin			Authenticate Database with Account		
* Datab	base Account:	100					
* Databa	ise Password:	**********		Ф	Test Connectivity		
Destination Database							
* Ir	nstance Type:	MongoDB Instance		•			
* Inst	tance Region:	China (Hangzhou)		•			
* MongoDB	Instance ID:	dds-bt		•			
* Dal	tabase Name:	admin			Authenticate Database with Account		
* Datab	base Account:						
* Databa	ise Password:	•••••		Ø	Test Connectivity		
					Cancel Assess Data	Migration to Cloud	Set Whitelist and Next
Section	Paran	neter	Description				

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select ApsaraDB for MongoDB.
	Instance Region	Select the region where the source ApsaraDB for MongoDB instance resides. In this example, select <b>China (Qingdao)</b> .
	MongoDB Instance ID	Select the ID of the source ApsaraDB for MongoDB instance.
	Database	Enter the name of the authentication database. The database account is created in this database.
	Name	<b>Note</b> If the database account is root, enter admin.
Source		Enter the database account of the source ApsaraDB for MongoDB
Database	Dat abase Account	instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the database account.
	Database Password	Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the source database parameters based on the check results.
	Instance Type	Select MongoDB Instance.
	Instance Region	Select the region where the destination ApsaraDB for MongoDB instance resides. In this example, select <b>China (Hangzhou)</b> .
	MongoDB Instance ID	Select the ID of the destination ApsaraDB for MongoDB instance.
	Database	Enter the name of the authentication database. The database account is created in this database.
	Name	<b>Note</b> If the database account is root, enter admin.
Destinati on		

Betebase	Parameter	Description	
	Dat <i>a</i> base Account	Enter the database account of the destination ApsaraDB for MongoDB instance. For information about the permissions that are required for the account, see Permissions required for database accounts.	
	Dat abase Password	Enter the password of the database account. Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.	

#### ii. In the lower-right corner of the page, click Set Whitelist and Next.

**Note** DTS adds the CIDR blocks of DTS servers to the whitelists of the source and destination ApsaraDB for MongoDB instances. This ensures that DTS servers can connect to the source and destination ApsaraDB for MongoDB instances. After data migration is complete, you can remove the CIDR blocks of DTS servers from the whitelists. For more information, see Configure a whitelist for a sharded cluster instance.

#### iii. Select the migration types and the objects to be migrated.

Available   If you search globally, please expand If you search globally	Source and Destination     Configure Migration Types and Objer     Migration Types: Full Data Migration     Incremental Data Mi	gration	3.Map name modification	>	4.Pr	scheck	
Select All       Remove All         *Name batch change :       Yes         Information:       1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.         2. DDL operations are not supported during data migration because this can cause migration failures.	Available If you search globally, please expand Q	> <	Selected (To edit an object name and click Edit.) Learn more.	e or its filter	, hover over the object	Hover over ti required objection click Edit. In box that app modify the o of the destin database and columns to n	he ect and the dia ears, bject n. ation d select
<ul> <li>*Name batch          <ul> <li>No</li> <li>Yes change :</li> </ul> </li> <li>Information:         <ul> <li>Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.</li> <li>DDL operations are not supported during data migration because this can cause migration failures.</li> </ul> </li> </ul>	Select All						
	<ul> <li>*Name batch          <ul> <li>No</li> <li>Yes change :</li> </ul> </li> <li>Information:         <ul> <li>Data migration only copies the data and schema in the source databadata or schema in the source database.</li> <li>DDL operations are not supported during data migration because the source databadata or schema in the source database.</li> </ul> </li> </ul>	ase and saves is can cause mi	the copy in the destination databas	e. The proc	ess does not affect any		

Setting	Description
	<ul> <li>To perform only full data migration, select only Full Data Migration.</li> <li>To ensure service continuity during data migration, select Full Data Migration and Incremental Data Migration.</li> </ul>
Select the migrati on types	<ul> <li>Note</li> <li>DTS does not support incremental data migration for standalone instances.</li> <li>If Incremental Data Migration is not selected, we recommend that you do not write data to the source instance during full data migration. This ensures data consistency between the source and destination instances.</li> </ul>
	<ul> <li>a. Select one or more objects from the Available section.</li> <li>b. Click the &gt; icon to move the objects to the Selected section.</li> </ul>
Select the objects to be migrate d	<ul> <li>Note</li> <li>DTS cannot migrate data from the admin or local database.</li> <li>You can select databases, collections, or functions as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination instance, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
connect ions to the source or destinat ion databas e	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

iv. In the lower-right corner of the page, click **Precheck**.

? Note

- Before you can start the data migration task, DTS performs a precheck. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed

item to view details.

- After you troubleshoot the issues based on the causes, you can run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- v. After the data migration task passes the precheck, click Next.
- vi. In the **Confirm Settings** dialog box, configure the **Instance Class** parameter. Then, read and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- vii. Click Buy and Start to start the data migration task.
- 7. View the progress of the migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

- Incremental data migration
   An incremental data migration task does not automatically stop. You must manually stop the task.
  - i. Wait until **Incremental Data Migration** and **The data migration task is not delayed** appear in the progress bar of the data migration task. Then, stop writing data to the source instance for a few minutes. The latency of **incremental data migration** may be displayed in the progress bar.
  - ii. Wait until the state of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.
- 8. Switch your workloads to the destination ApsaraDB for MongoDB instance.

#### What to do next

If you no longer need the source instance, release the source instance.

- If the source instance is charged based on the pay-as-you-go billing method, you can release the instance. For more information, see Release an instance or a node.
- If the source instance is charged based on the subscription billing method, you are not allowed to release the instance.

## 8.12. Migrate data from a MaxCompute project to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a MaxCompute project to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS).

#### Prerequisites

The following operations are performed:

- Activate MaxCompute
- Create a project

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- DTS automatically creates a database in the destination ApsaraDB RDS for MySQL instance. However, if the name of the source database is invalid, you must manually create a database in the destination instance before you configure the data migration task.

(?) Note For more information about the naming conventions of ApsaraDB RDS and how to create a database, see Create databases and accounts for an ApsaraDB RDS for MySQL instance.

- To ensure data consistency, we recommend that you do not write data to the source MaxCompute project during full data migration.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination instance, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination instance after the task is resumed.
- MaxCompute does not support the PRIMARY KEY constraint. If network errors occur, DTS retries the data migration task. In this case, duplicate data records may be migrated to the destination tables that do not have primary keys.
- DTS cannot migrate incremental data from a MaxCompute project to an ApsaraDB RDS for MySQL instance.

**Note** To ensure data consistency, we recommend that you do not write data to the source MaxCompute project during data migration.

#### Supported destination database types

You can use DTS to migrate data to the following types of MySQL databases:

- Self-managed database hosted on Elastic Compute Service (ECS)
- Self-managed database connected over Express Connect, VPN Gateway, or Smart Access Gateway
- Self-managed database connected over Database Gateway
- ApsaraDB RDS for MySQL instances that are owned by the same Alibaba Cloud account as the MaxCompute project or a different Alibaba Cloud account from the MaxCompute project

In this topic, an **ApsaraDB RDS for MySQL instance** is used to describe how to configure a data migration task. You can also follow the procedure to configure data migration tasks for other types of MySQL databases.

**Note** If your destination database is a self-managed MySQL database, you must deploy the network environment for the source database. For more information, see **Preparation overview**.

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see Pricing.

#### **Migration types**

Migration type	Description			
	DTS migrates the schemas of required objects to the destination database. DTS supports schema migration for views, tables, and databases.			
Schema migration	• Warning In this topic, the source and the destination databases are heterogeneous databases. DTS does not ensure that the schemas of the source and destination databases are consistent after schema migration. We recommend that you evaluate the impact of data type conversion on your business. For more information, see Data type mappings between heterogeneous databases.			
	DTS migrates the historical data of required objects to the destination database.			
Full data migration	<b>Note</b> During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.			

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.

4. Configure the source and destination databases.

1.Configure Source and Destination	2.Authorize MaxCompute Account > 3	.Configure	Migration Types and	$\rangle$	4.Precheck
* Task Name:	axCompute to RDS				
		time it is set	encounted to use the data.	numericantian from	ation which has better actions
stability and richer functions. Click to buy data	a synchronization Click to see the difference between data migration	ration and syr	commended to use the data : ichronization	synchronization fun	ction, which has better network
Source Database					
* Instance Type:	MaxCompute	~	DTS support type		
* Instance Region:	China (Hangzhou)	~			
# Dreight					
Project:	dtstest123				
Destination Database					
* Instance Type:	RDS Instance	~			
* Instance Region					
Instance Region:	China (Hangzhou)	~			
* RDS Instance ID:	rm a second and	-			
* Database Account:	dtetect				
balababerrecountr					
* Database Password:	•••••	<b>∢</b> >	Test Connectivity		
* Encryption:	Non-encrypted      SSL-encrypted				
				Cancel	Set Whitelist and Next

Section	Parameter	Description			
	Instance Type	Select MaxCompute.			
	Instance Region	The region where the source MaxCompute project resides.			
Source Database	Project	The name of the MaxCompute project. You can search for a project on the Workspaces page in the DataWorks console.			
	Instance Type	Select RDS Instance.			
	InstanceThe region where the destination ApsaraDB RDS for IRegioninstance resides.				
	RDS Instance ID	The ID of the destination ApsaraDB RDS for MySQL instance.			

Section	Parameter	Description		
Destination	Dat abase Account	The database account of the destination ApsaraDB RDS for MySQL instance. The account must have read and write permissions on the destination database.		
Database		The password of the database account. After you specify the destination database parameters, click <b>Test</b> <b>Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid.		
Dat abase Password		<b>Note</b> If the specified parameters are valid, the <b>Passed</b> message is displayed. If the <b>Failed</b> message is displayed, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.		
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the ApsaraDB RDS for MySQL instance before you configure the data migration task. For more information, see <b>Configure SSL</b> <b>encryption for an ApsaraDB RDS for MySQL instance</b> .		

#### 5. In the lower-right corner of the page, click Set Whitelist and Next.

#### □ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 6. In the lower-right corner of the page, click **Next**. In this step, the permissions on the MaxCompute project are granted to the migration account. The following figure provides an example.

1.Configure Source and Destination	2.Authorize MaxCompute Account	3.Configure Migration Types and	A.Precheck
To synchronize data to a MaxCompu	te instance, you must grant the following permissio	ns of project dtstest123 to the synchronization	account.
CreateTable			
CreateInstance			
CreateResource			
CreateJob			
List			
			Cancel Previous

7. Select the migration type and the objects to be migrated.

1.Configure Source and	2.Authorize MaxCompute Account 3.0	Configure Migration Types and	4.Precheck
Migration Types: Schema Migration During full data migration, data update For data consistency, we recommend When the CPU of the target instance M	on Full Data Migration es in the source database are not migrated to the de that you select Schema Migration, Full Data Migration MySQL is lower than 2 cores, it is recommended to ch	stination instance. η, and Incremental Data Migration. ose the slow log, <b>View More</b>	
Available Expand the tree before you perform a distest123 Tables	aglo Q	Selected (To edit an object name or its filts Edit.) Learn more.	er, hover over the object and dick
		Remove All	
*Rename Databases and Tables: * Retry Time for Failed Connection *Source table DMS_ONLINE_ Do you want to copy the temporary table to the target database during DDL:	Do Not Change Database and Table Names     720     Minutes     ?     Yes     No     ?	<ul> <li>Change Database and Table Names</li> </ul>	
*Target library object name case policy:	DTS default policy      Consistent with source     Consistent with the target library default po     Consistent with the target library default po	e library licy(Capitalize) licy(a lowercase letter)	
Information: 1. Data migration only copies the data an in the source database. 2. Do not do DDL operation during struct	d schema in the source database and saves the copy ure and full migration, otherwise the task may fail	y in the destination database. The process does	not affect any data or schema
		Cancel Previous	Save Precheck
#### Data Migration Migrate data betwe en instances of the same Alibaba Cl oud account

Select Schema Migration and Full Data Migration.         Image: Select Schema Migratin.         Image:	) an
the migratio n types       ? Note DTS cannot migrate incremental data from a MaxCompute project to ApsaraDB RDS for MySQL instance.         Select one or more objects from the Source Objects section and click the > icor the objects to the Selected Objects section.	o an to add
Select one or more objects from the <b>Source Objects</b> section and click the <b>Selected Objects</b> section.	to add
the objects to the Selected Objects section.	
<ul> <li>Select the objects that you want to migrate</li> <li>By default, after an object is migrated to the destination database, the n of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database, the n of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database, the n of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>	ame Iase.
Specify the retry timeBy default, if DTS fails to connect to the source or destination database, DTS retries the following 12 hours. You can specify the retry time range based on your business for failed connectiBy default, if DTS fails to connect to the source or destination database, DTS retries 	s within s the gration
ons to         the       ⑦ Note       When DTS retries a connection, you are charged for the DTS instance         source       recommend that you specify the retry time range based on your business needs.         or       can also release the DTS instance at your earliest opportunity after the source and destination instances are released.         on       databas	. We You d

8. In the lower-right corner of the page, click **Precheck**.

#### ➡ Notice

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the next to each failed item

to view details.

- After you troubleshoot the issues based on the causes, you can run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click **Next**.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.

(?) Note We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

12. Switch your workloads to the ApsaraDB RDS for MySQL instance.

#### What's next

The members of the source MaxCompute project have read permissions. The database account of the destination ApsaraDB RDS for MySQL instance has read and write permissions. To ensure database security, perform the following operations after data migration is complete: 1. Remove the members of the source MaxCompute project or modify their roles. 2. Delete the database account of the destination ApsaraDB RDS for MySQL instance. For more information, see Add project members and configure roles and Delete accounts.

# 9.Migrate data between instances of different Alibaba Cloud accounts 9.1. Migrate data between PolarDB for MySQL clusters of different Alibaba Cloud accounts

PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. PolarDB is compatible with the MySQL database engine and features high performance, high availability, high reliability, and ease of use. This topic describes how to migrate data between PolarDB for MySQL clusters that are owned by different Alibaba Cloud accounts by using Data Transmission Service (DTS).

#### Prerequisites

- The source and destination PolarDB for MySQL clusters are created. For more information, see Create a PolarDB for MySQL cluster.
- The binary logging feature is enabled for the source PolarDB for MySQL cluster. For more information, see Enable binary logging.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination database after the task is resumed.

#### **Migration types**

DTS supports schema migration, full data migration, and incremental data migration. For more information, see Terms.

**Note** When you migrate data between PolarDB for MySQL clusters, you can select all of the supported migration types to ensure service continuity.

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

#### Permissions required for database accounts

Database	Required permissions
Source PolarDB for MySQL cluster	The read permission on the objects to be migrated
Destination PolarDB for MySQL cluster	The read and write permissions on the objects to be migrated

Onte For more information about how to create and authorize a database account, see Create a database account.

#### Before you begin

- 1. Log on to the Alibaba Cloud Management Console by using the Alibaba Cloud account that owns the source PolarDB for MySQL cluster.
- 2. Create a RAM role and authorize this role to access the cloud resources of the Alibaba Cloud account that owns the source PolarDB for MySQL cluster. For more information, see Configure RAM authorization for data migration or synchronization from a self-managed database in a VPC across different Alibaba Cloud accounts.

**Note** To migrate data between PolarDB for MySQL clusters of different Alibaba Cloud accounts, you can specify the source PolarDB for MySQL cluster as a self-managed database connected over Express Connect. You can also request a public endpoint for the source PolarDB for MySQL cluster and specify the source cluster as a self-managed database with a public IP address. In this case, you do not need to create a RAM role and configure authorization.

#### Procedure

- 1. Log on to the DTS console by using the Alibaba Cloud account that owns the destination PolarDB for MySQL cluster.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Select User-Created Database Connected over Express Connect, VPN Gateway, or Smart Access Gateway as the instance type. Then, click VPC of Another Alibaba Cloud Account next to the Peer VPC field.

Note To migrate data between PolarDB for MySQL clusters of different Alibaba Cloud accounts, you can specify the source PolarDB for MySQL cluster as a self-managed database connected over Express Connect. You can also request a public endpoint for the source PolarDB for MySQL cluster and specify the source cluster as a self-managed database with a public IP address.

1.Configure Source and Destination	2.Configure Migration Types and Objects	3.Adva	nced Settings	$\rightarrow$	4.Precheck
* Task Name: 🔒	l Mildae				
Source Database					
* Instance Type:	User-Created Database with Public IP Address	• DT	S support type		
* Instance Region:	China (Hangzhou)	• Get	IP Address Segment of	DTS	
* Database Type:	MySQL	•			
Hostname or IP Address:					
* Port Number:	3306				
* Database Account:					
* Database Password:		Þ	Test Connectivity		

6. Configure the source and destination databases.

#### Dat a Transmission Service

#### Data Migration Migrate data betwe en instances of different Alibaba Cl oud accounts

* Tas	k Name: P	olarDB MySQL		
Source Database				
* Inst	ance Type:	User-Created Database (	Connected Over Express Connect, VPN 🔻	DTS support type
* Instan	nce Region:	Singapore	•	Guide
*Apsara Stack Tenant A	Account ID:			
*r	Role Name:	e: ram-for-dts		Authorize Role Across Accounts
+	Peer VPC:	C: vpc		Proprietary network of the current login account
* Data	base Type:	pe: MySQL V		
+ 1	IP Address:	172.168		
* Po	rt Number:	3306		-
* Databas	se Account:	dtstest		
* Database	Password:	•••••	<i>ه</i> ›	Test Connectivity 📀 Passed
Destination Database				
* Inst	ance Type:	PolarDB	v	
* Instan	nce Region:	Singapore	v	
* PolarDB In	nstance ID:			
* Databas	se Account:	dtstest		
* Database	Password:	*****	<b>4</b> >	Test Connectivity 📀 Passed
				Cancel Set Whitelist and Next
Section	on Parameter Descr		Description	
N/A	DTS automatically generates a task name. We recommend that youTask Namespecify an informative name for easy identification. You do not needto use a unique task name.		generates a task name. We recommend that you tive name for easy identification. You do not need k name.	
	Instance Type Select User-Cr Connect, VPN		Select User-Creat Connect, VPN Ga	ed Database Connected over Express teway, or Smart Access Gateway.
	Instance Region		Select the region where the source PolarDB cluster resides.	
			Enter the ID of the PolarDB cluster.	Alibaba Cloud account that owns the source
	Alibaba Cloud Account ID Role Name		Note To a owns the source Management cor displayed on the	obtain the ID of the Alibaba Cloud account that PolarDB cluster, you must log on to the Account nsole by using this account. The account ID is Security Settings page.
			Account Management Back Information Security Settings Resinante Verification Contact Management [2]	Center         2 Basic Information           Logon Account:         dor <sup>um</sup> @Hest aligund com Edit         Verlied:         No real-same authentication         Go Is centeration           Account ID         dor <sup>um</sup> @Hest aligund com Edit         Kegistration time:         2021-05-25 11.39 00         Centeration           Three party binding         No binding Detail         Verlied:         Verlied:         2021-05-25 11.39 00
			Enter the name of begin.	the RAM role that you created earlier in <mark>Before you</mark>

Section	Parameter	Description				
		Select the ID of the VPC where the source PolarDB cluster resides. To obtain the VPC ID, you must log on to the PolarDB console by using the Alibaba Cloud account that owns the source PolarDB cluster. On the Clusters page, click the ID of the source PolarDB cluster. The VPC ID is displayed in the <b>Basic Information</b> section.				
Source Dat abase	Peer VPC	Cluster         Running         Leg On to Dilutoure         Migrate Bom Other Database         Cone Cluster           Overview         Basic Information         Cluster ID         Cluster ID         Cluster ID         Cluster Name         Edit           Obigoostics and Optime         Diagnostics and Optime         Singapore         Zones         Singapore Zone A (Primary), Singapore           VPC         vpc         VSwitch         VSwitch         VSwitch				
	Database Type	Select MySQL.				
	IP Address	Enter the private IP address of the source PolarDB cluster. In this example, enter <b>172.16.20.20</b> . You can obtain the private IP address by pinging the <b>VPC-facing endpoint</b> of the source PolarDB cluster.				
	Port Number	Enter the service port number of the source PolarDB cluster. The default port number is <b>3306</b> .				
	Dat abase Account	Enter the database account of the source PolarDB cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.				
	Dat abase Password	Enter the password of the database account. <b>Note</b> After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.				
	Instance Type	Select <b>PolarDB</b> .				
	Instance Region	Select the region where the destination PolarDB cluster resides.				
	PolarDB Instance ID	Select the ID of the destination PolarDB cluster.				
	Dat abase Account	Enter the database account of the destination PolarDB cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.				
Destinatio n Database						

Section	Parameter	Description
		Enter the password of the database account.
	Database Password	<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.

#### 7. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note DTS adds the CIDR blocks of DTS servers to the whitelists of the source and destination PolarDB clusters. This ensures that DTS servers can connect to the source and destination PolarDB clusters.

8. Select the migration types and the objects to be migrated.

1.Configu	re Source and	2.Configure Migration Ty	/pes and	3.Advanced Settings	>	4.Precheck
<ul> <li>Migration T triggers. For m</li> </ul>	ypes: 🗹 Schema Migratio ore information, see Refer	on 🔽 Full Data Migration	☑ Incremental [	Data Migration Note: Incr	emental data migra	ation does not support
Note: do not cleans up th	t clean up the incremental e log too early, the DTS in	data log generated by the sour cremental task may fail	ce database after the	DTS task is started when the	DTS full task is rur	nning. If the source database
Data migrati between Ap: For long-ten	ion applies to short-term m sara Stack databases. m data synchronization in i	igration scenarios. Typical scen real time, use the data synchror	arios include migratin nization feature.	g data to the doud, scaling a	nd sharding databa	ises, and migrating data
Available				Selected (To edit an object Edit.) Learn more.	t name or its filter,	, hover over the object and click
Expand the	tree before you perform a	glol   Q				Q
🗆 📑 dtst	estdata			📼 dtstestdata (20	)biects)	
	adiés iews			customer		
	IEWS		> <	i order		
Select All *Rename Data	bases and Tables:	Do Not Change Databas	e and Table Names	Remove All <ul> <li>Change Database and</li> </ul>	Table Names	
* Retry Time fo	or Failed Connection	720 Minute	5 🕜			
*Source table I want to copy th the target data	DMS_ONLINE_Do you ne temporary table to base during DDL:	🔾 Yes 💿 No 🕐				
Information: 1. Data migrati in the source d 2. Do not do D	on only copies the data an atabase. DL operation during structi	d schema in the source databas ure and full migration, otherwise	e and saves the copy e the task may fail	in the destination database. "	The process does r	not affect any data or schema
				Cancel	Previous	Save Precheck
etting	Description					
	<ul> <li>To perform</li> </ul>	n only full migratic	on, select <b>Sc</b> l	hema Migration	and Full D	Data Migration.
elect	<ul> <li>To ensure</li> <li>Data Migr</li> </ul>	service continuity o ation, and Incren	during data r n <b>ental Dat</b> a	nigration, select S Migration.	Schema M	igration, Full
the migratio n types Notice If Incremental Data Migration is not selected, we recommend that you do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.						

#### Dat a Transmission Service

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the <b>&gt;</b> icon to move the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Notice</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
the source or destinati on databas	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
е	

Setting	Description
Specify whether to copy tempora ry tables to the destinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.
on databas e when DMS	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.
perform s online DDI	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
operatio ns on the source	<b>Note</b> If you select No, the tables in the destination database may be locked.
table	

#### 9. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the 🕧 icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.

#### 10. After the task passes the precheck, click Next.

- 11. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 12. Click Buy and Start to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



#### What's next

The database accounts that are used for data migration have the read and write permissions. After data migration is complete, you must delete the database accounts to ensure security.

# 9.2. Migrate data between ApsaraDB for MongoDB instances of different Alibaba Cloud accounts

This topic describes how to migrate data between ApsaraDB for MongoDB instances of different Alibaba Cloud accounts by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration. When you migrate data between ApsaraDB for MongoDB instances, you can select both of the supported migration types to ensure service continuity.

#### Prerequisites

• The source instance is a standalone instance or a replica set instance. If the source instance is a sharded cluster instance, we recommend that you use the built-in tools of MongoDB to migrate data. For more information, see Migrate a self-managed MongoDB database to ApsaraDB for MongoDB by using tools provided by MongoDB.

**?** Note DTS does not support incremental data migration from a standalone instance. For more information, see Migration types.

• The destination instance is created in the destination region. For more information, see Create a standalone instance, Create a replica set instance, or Create a sharded cluster instance.

**?** Note The available storage space of the destination instance must be larger than the total size of the data in the source instance.

#### Precautions

- DTS consumes the resources of the source and destination instances during full data migration. This may increase the loads of the database servers. If you migrate a large volume of data or the server specifications cannot meet your requirements, database services may become unavailable. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- DTS does not support incremental data migration from a standalone instance. To ensure data consistency, do not write data to the source instance during full data migration.
- If the source and destination MongoDB databases use different versions or storage engines, make sure that your applications can run on both databases. For more information about the versions and storage engines that are supported by ApsaraDB for MongoDB, see MongoDB versions and storage engines.

•

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba
Incremental data migration	Charged. For more information, see Pricing.	Cloud over the Internet. For more information, see Pricing.

#### **Migration types**

Migration type	Description
Full data migration	DTS migrates the historical data of required objects from the source MongoDB database to the destination MongoDB database.
	Note The following types of objects are supported: database, collection, and index.
	After full data migration is complete. DTS synchronizes incremental data from the
	source MongoDB database to the destination MongoDB database.
Incremental data migration	<ul> <li>Note</li> <li>The create and delete operations that are performed on databases, collections, and indexes can be synchronized.</li> <li>The create, delete, and update operations that are performed on documents can be synchronized.</li> </ul>

#### Permissions required for database accounts

Database	Full data migration	Incremental data migration
Source ApsaraDB for MongoDB instance	The read permissions on the source database	The read permissions on the source database, the admin database, and the local database
Destination ApsaraDB for MongoDB instance	The read and write permissions on the destination database	The read and write permissions on the destination database

**Note** For more information about how to create and authorize a database account, see Manage user permissions on MongoDB databases.

#### Before you begin

- 1. Log on to the ApsaraDB for MongoDB console by using the Alibaba Cloud account that owns the source instance.
- 2. Apply for a public endpoint for the source instance. For more information, see Apply for a public endpoint.
- 3. Add the CIDR blocks of DTS servers to the whitelist of the source instance. For more information, see Configure a whitelist or an ECS security group for an ApsaraDB for MongoDB instance.

(?) Note You can determine the CIDR blocks you need to add based on the region where the destination instance resides. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.

For example, if the source instance resides in the China (Hangzhou) region and the destination instance resides in the China (Shenzhen) region, you must add the CIDR blocks of DTS servers in the China (Shenzhen) region to the whitelist of the source instance.

#### Procedure

1.

2.

3. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.

4.

- 5. In the upper-right corner of the page that appears, click Migrate Databases.
- 6. On the Create Migration Task page, create a data migration task.
  - i. Configure the source and destination databases.

1.Configure Source and Destination	Databases 2.Configure Migration Types and Objects	> 3.Map name	modification	$\rangle$	4.Precheck
* Task Name: d					
Course Database					
Source Database					
* Instance Type:	User-Created Database with Public IP Address				
* Instance Region:	China (Shanghai) 🔻	Get IP Address Segment of	f DTS		
* Database Type:	MongoDB 🔻				
* Hostname or IP Address:	dds-bp1mongodb.rds.aliyuncs.co				
* Port Number:	3717				
Database Name:	admin	Authenticate Database with	Account		
Database Account:					
Database Password:	4	Test Constantiation			
	4/	Test Connectivity			
	······	Test Connectivity			
Destination Database	······	Test Connectivity			
Destination Database • Instance Type:	MongoDB Instance	Test Connectivity			
Destination Database  Instance Type: Instance Region:	MongoDB Instance v China (Hangzhou) v	Test Connectivity			
Destination Database  Instance Type: Instance Region: MongoDB Instance ID:	MongoDB Instance   China (Hangzhou)  dds-br	Test Connectivity			
Destination Database  Instance Type: Instance Region: MongoDB Instance ID: Database Name:	MongoDB Instance  China (Hangzhou) dds-by admin	Authenticate Database with	Account		
Destination Database  Instance Type: Instance Region: MongoDB Instance ID: Database Name: Database Account:	MongoDB Instance   China (Hangzhou)  dds-bi admin	Authenticate Database with	Account		
Destination Database  Instance Type: Instance Region: MongoDB Instance ID: Database Name: Database Account: Database Password:	MongoDB Instance   China (Hangzhou)  dds-br admin	Authenticate Database with	Account		
Destination Database  Instance Type: Instance Region: MongoDB Instance ID: Database Name: Database Account: Database Password:	MongoDB Instance   China (Hangzhou)  dds-bt admin	Authenticate Database with	Account		
Destination Database  Instance Type: Instance Region: MongoDB Instance ID: Database Name: Database Account: Database Password:	MongoDB Instance   China (Hangzhou)  dds-bi admin	Authenticate Database with Test Connectivity	Account		

Section	Parameter	Description				
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.				
	Instance Type	Select User-Created Database with Public IP Address.				
	Instance Region	Select the region where the source instance resides.				
	Database Type	Select MongoDB.				
	Hostname or IP Address	Enter the domain name obtained from the public endpoint of the source instance. For example, enter dds-1udxxxxxxx-pub.mongodb.rds.aliyuncs.com.				
	Port Number	Enter <b>3717</b> , which is the service port of the source instance.				
	Database	Enter the name of the authentication database. The database account is created in this database.				
	Name	<b>Note</b> If the database account is root, enter admin.				
	Dat abase Account	Enter the database account of the source ApsaraDB for MongoDB instance. For information about the permissions that are required for the account, see Permissions required for database accounts.				

Section	Parameter	Description
Database		Enter the password of the database account.
	Dat abase Password	Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the source database parameters based on the check results.
		This parameter specifies whether to use an encrypted connection between DTS and the source instance. Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you select <b>SSL-encrypted</b> , you must enable SSL encryption for the source instance before you configure the data migration task.
	Encryption	<ul> <li>Note</li> <li>You can select SSL-encrypted only for MongoDB Atlas databases.</li> <li>If you select SSL-encrypted, more CPU resources will be consumed.</li> </ul>
	Instance Type	Select MongoDB Instance.
	Instance Region	Select the region where the destination ApsaraDB for MongoDB instance resides.
	MongoDB Instance ID	Select the ID of the destination ApsaraDB for MongoDB instance.
Destinati	Database Name	Enter the name of the authentication database. The database account is created in this database.           Image: The database           Image: Note         Image: The database account is root, enter admin.
	Database Account	Enter the database account of the destination ApsaraDB for MongoDB instance. For information about the permissions that are required for the account, see Permissions required for database accounts.
on Database		

Section	Parameter	Description		
		Enter the password of the database account.		
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.		

#### ii. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note After data is migrated, you can remove the CIDR blocks of DTS servers from the whitelist of the source instance if they are no longer needed. For more information, see Configure a whitelist for an ApsaraDB for MongoDB instance.

#### iii. Select the migration types and the objects to be migrated.

1.Configure Source and Destination	2.Configure Migration Types and Obje	cts			heck
<ul> <li>Migration Types:  Full</li> </ul>	Data Migration 🛛 🗹 Incremental Data Mi	gration			
Available If you search globally, pl If 🖕 admin	lease expand	<b>&gt;</b> <	Selected (To edit an object name o and click Edit.) Learn more.	or its filter, hover over the object	Hover over the required object and click Edit. In the dialog box that appears, modify the object name of the destination database and select the columns to migrate.
Select All			Remove All		
*Name batch @ change : Information: 1. Data migration only copie data or schema in the sour 2. DDL operations are not s	No Yes es the data and schema in the source database. upported during data migration because this	ase and saves	the copy in the destination database,	The process does not affect any	
			C	Cancel Previous Save	Precheck
Setting	Description				

Setting	Description			
Select the migration types	<ul> <li>Description</li> <li>To ensure data consistency, we recommend that you select both Full Data Migration and Incremental Data Migration.</li> <li>If you select only Full Data Migration, the data updated during the migration process may not be migrated from the source instance to the destination instance.</li> <li>If you select only Incremental Data Migration, take note of the following information:</li> <li>After the migration task is started, DTS migrates only incremental data generated in the source instance.</li> <li>Triggers cannot be synchronized during incremental data migration. For more information, see Configure a data synchronization task for source database that contains a trigger.</li> <li>Note</li> <li>DTS does not support incremental data migration for standalone instances.</li> <li>If Incremental Data Migration is not selected, we recommend that you do not write data to the source instance during full data migration. This ensures data consistency between the source and destination instances.</li> </ul>			
	<ul> <li>a. Select one or more objects from the Available section.</li> <li>b. Click the &gt; icon to move the objects to the Selected section.</li> </ul>			
Select the objects to be migrated	<ul> <li>a. Select one or more objects from the Available section.</li> <li>b. Click the  icon to move the objects to the Selected section.</li> <li>⑦ Note <ul> <li>DTS cannot migrate data from the admin or local database.</li> <li>You can select databases, collections, or functions as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination instance, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.</li> </ul> </li> </ul>			

Setting	Description			
Specify the retry time for failed connections to the source or destination database	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.			
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.			

#### iv. In the lower-right corner of the page, click Precheck.

#### ? Note

- Before you can start the data migration task, DTS performs a precheck. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed

item to view details.

- After you troubleshoot the issues based on the causes, you can run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- v. After the data migration task passes the precheck, click Next.
- vi. In the **Confirm Settings** dialog box, configure the **Instance Class** parameter. Then, read and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- vii. Click Buy and Start to start the data migration task.
- 7. View the progress of the migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

- i. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the source instance for a few minutes. The latency of incremental data migration may be displayed in the progress bar.
- ii. Wait until the state of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.

8. Switch your workloads to the destination ApsaraDB for MongoDB instance.

#### What to do next

If you no longer need the source instance, release the source instance.

- If the source instance is charged based on the pay-as-you-go billing method, you can release the instance. For more information, see Release an instance or a node.
- If the source instance is charged based on the subscription billing method, you are not allowed to release the instance.

#### References

If you migrate data to a sharded cluster instance, you can configure data sharding based on your needs. For more information, see Configure sharding to maximize the performance of shards.

# 9.3. Migrate data between RDS instances of different Alibaba Cloud accounts

This topic describes how to migrate data between RDS instances of different Alibaba Cloud accounts by using Data Transmission Service (DTS).

#### Prerequisites

The available storage space of the destination instance is larger than the total size of the data in the source instance.

#### Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

#### Permissions required for database accounts

Instance	Schema migration	Full data migration	Incremental data migration
Source RDS instance	The read and write permissions	The read and write permissions	The read and write permissions
Destination RDS instance	The read and write permissions	The read and write permissions	The read and write permissions

#### Before you begin

Log on to the Resource Access Management (RAM) console by using the Alibaba Cloud account that owns the source instance (Account A). Specify the Alibaba Cloud account that owns the destination instance (Account B) as a trusted account. Then, authorize Account B to access the cloud resources of Account A by using DTS. For more information, see Configure RAM authorization for cross-account data migration and synchronization.

#### Procedure

- 1. Log on to the DTS console by using the Alibaba Cloud account that owns the destination RDS instance.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the **Migration Tasks** page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. In the Source Database section, set Instance Type to RDS Instance, and click RDS Instances of Other Apsara Stack Accounts next to the RDS Instance ID field.

1.Configure Source and	2.Configure Migration Typ	es and $>$	3.Advanced Settings	<u> </u>	I.Precheck
* Task Name: RI	DS_TO_RDS				
Note which has better network stability and richer fr	: If you need to carry out incre unctions. Click to buy data syn	emental data migrati chronization Click to	on for a long time, it is rec see the difference betwe	commended to use the data en data migration and syncl	synchronization function, pronization
Source Database					
* Instance Type	RDS Instance		~	DTS support type	
* Instance Region:	China (Hangzhou)		~		
* RDS Instance ID;			-	RD5 Instances of Other A	psara Stack Accounts
* Database Account:					
* Database Password:			Þ	Test Connectivity	
Destination Database					
* Instance Type:	RDS Instance		~		
* Instance Region:	China (Hangzhou)		~		
* RDS Instance ID:	Select an RDS instance.		-		
* Database Account:					
* Database Password:			4>	Test Connectivity	
			Cancel Set White	list and Next Experier	ce the new console

6. Configure the source and destination databases.

1.Configure Source and Destination	2.Configure Migration Types and $>$ 3.4	Advanced Setti	ings 🔰 4.Precheck
_			
* Task Name: F	DS_TO_RDS		
Not has better network stability and richer function	e: If you need to carry out incremental data migration for a lo ns. Click to buy data synchronization Click to see the differe	ong time, it is rea nce between dat	commended to use the data synchronization function, which a migration and synchronization
Curry Database			
Source Database			
* Instance Type:	RD5 Instance	~	DTS support type
* Instance Region:	China (Hangzhou)	~	
*Apsara Stack Tenant Account ID of RDS			
Instance			Guide
*Role Name:	- and a final sector of the se		Authorize Role Across Accounts
*RDS Instance ID:		-	RDS Instances of Current Account
* Database Account:			
* Database Password:		4>	Test Connectivity
Destination Database			
* Instance Tunes	DDC Taskasa		
Instance Type:	RUS Instance	•	
* Instance Region:	China (Hangzhou)	~	
* RDS Instance ID:	rm-	•	
* Database Account:			
* Database Password:	•••••	<⊅	Test Connectivity
* Encryption:	Non-encrypted O SSL-encrypted		
		Cancel Se	t Whitelist and Next Experience the new console

Parameter	Description
Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to specify a unique task name.

Parameter	Description				
	• Instance Type: Select <b>RDS Instance</b> .				
	<b>Note</b> If you have selected <b>RDS Instance</b> as the instance type, you can skip this step.				
	• Instance Region: Select the region where the source RDS instance resides.				
	ONDE You can select different regions for the source and destination RDS instances.				
	<ul> <li>Apsara Stack Tenant Account ID of RDS Instance: Enter the ID of the Alibaba Cloud account that owns the source instance.</li> </ul>				
Source Database	<b>Note</b> To obtain the ID of the Alibaba Cloud account that owns the source instance, you must log on to the Account Management console by using this account. The account ID is displayed on the Security Settings page.				
	<ul> <li>Role Name: Enter the role name configured for the Alibaba Cloud account that owns the source instance. For more information, see Configure RAM authorization for cross-account data migration and synchronization.</li> <li>RDS Instance ID: Select the ID of the source RDS instance.</li> </ul>				
	<b>Note</b> If an alert message appears when you select an RDS instance ID, modify the parameter values as prompted. For more information, see FAQ.				
	<ul> <li>Database Account: Enter the database account of the source RDS instance. For information about the permissions that are required for the account, see Permissions required for database accounts.</li> <li>Database Password: Enter the password of the database account.</li> </ul>				

Parameter	Description			
	• Instance Type: Select <b>RDS Instance</b> .			
	Note You can select different regions for the source and destination RDS instances.			
	• Instance Region: Select the region where the destination RDS instance resides.			
	RDS instance ID: Select the ID of the destination RDS instance.			
Destination Database	<ul> <li>Database Account: Enter the database account of the destination RDS instance.</li> <li>For information about the permissions that are required for the account, see</li> <li>Permissions required for database accounts.</li> </ul>			
	• Database Password: Enter the password of the database account.			
	<ul> <li>Encryption: Select Non-encrypted or SSL-encrypted. In this example, select Non-encrypted.</li> </ul>			
	<b>Note</b> If you want to select <b>SSL-encrypted</b> , you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see <b>Configure SSL encryption for an ApsaraDB RDS for MySQL instance.</b>			

#### 7. In the lower-right corner of the page, click Set Whitelist and Next.

#### 🗘 Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 8. Select the objects to be migrated and the migration types.

1.Configu	ure Source and	2.Configure Migration Ty	rpes and	3.Advanced Settings	$\rightarrow$	4.Precheck
* Migration 1 triggers. For m	Types: 🔽 Schema Migrat	ion 🔽 Full Data Migration rence	✓ Incremental [	Data Migration Note: Increm	mental data mig	ration does not support
Note: do no cleans up th	t clean up the incrementa ie log too early, the DTS i	l data log generated by the sourc ncremental task may fail	e database after the	DTS task is started when the D	TS full task is ru	inning. If the source database
Data migrat between Ap For long-ter	ion applies to short-term sara Stack databases, m data synchronization in	migration scenarios. Typical scena real time, use the data synchron	arios include migratir ization feature.	g data to the cloud, scaling and	l sharding datab	ases, and migrating data
Available				Selected (To edit an object Edit.) Learn more.	name or its filte	r, hover over the object and dic
Expand the	e tree before you perform	a glol 🛛 🔍				Q
🗆 💼 dtst	estdata			🚍 dtstestdata (201	viects)	
	ables lieure			customer		
÷ 📻 V	liews		> <	i order		
Select All	bases and Tables:	De Net Change Debuter	and Table Names	Remove All	bla Namar	
*Rename Data	abases and Tables:	Do Not Change Database	and Table Names	Change Database and I.	able Names	
*Source table want to copy t the target data	DMS_ ONLINE_ Do you he temporary table to abase during DDL:	Ves No 🤇				
1. Data migrat in the source o 2. Do not do D	ion only copies the data a latabase. IDL operation during struc	nd schema in the source databasi ture and full migration, otherwise	e and saves the copy the task may fail	in the destination database. Th	e process does	not affect any data or schema
				Cancel	Previous	Save Precheck
etting	Description					
	• To perfor Migratior	m only full data miç 1.	gration, sele	ct Schema Migra	t <b>ion</b> and	Full Data
elect	<ul> <li>To ensure service continuity during data migration, select Schema Migration, Full</li> <li>Data Migration, and Incremental Data Migration.</li> </ul>					
igratio types	? Note	If Incremental I	Data Migrat	ion is not selecte	d, do not	write data to the

the source and destination databases.

#### Dat a Transmission Service

Setting	Description
	Select one or more objects from the <b>Available</b> section and click the > icon to add the objects to the <b>Selected</b> section.
Select the objects to be migrate d	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated. If you select tables or columns as the objects to be migrated, DTS does not migrate other objects such as views, triggers, and stored procedures to the destination database.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination RDS instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS reconnects to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

Setting	Description				
Specify whether to copy tempora ry tables to the destinati on databas e when DMS perform s online DDL operatio ns on the source table	lf da op	you use Data Management (DMS) to perform online DD atabase, you can specify whether to migrate temporar perations. Yes: DTS migrates the data of temporary tables gene	L operations on the source y tables generated by online DDL erated by online DDL operations.		
		<b>Note</b> If online DDL operations generate a large migration task may be delayed.	ge amount of data, the		
	o	<b>No</b> : DTS does not migrate the data of temporary tab operations. Only the original DDL data of the source d	les generated by online DDL atabase is migrated.		
		Note If you select No, the tables in the desti	nation database may be locked.		

#### 9. Click Precheck.

? Note

- A precheck is performed before the migration task starts. The migration task only starts after the precheck succeeds.
- If the precheck fails, click the

icon next to each failed check item to view the related details. Fix the issues as instructed and run the precheck again.

- 10. After the data migration task passes the precheck, click Next.
- 11. In the **Confirm Settings** dialog box, configure the **Channel Specification** parameter. Then, read and select **Data Transmission Service (Pay-as-you-go) Service Terms**.
- 12. Click Buy and Start to start the data migration task.
  - Full data migration

Do not manually stop a full data migration task. If you manually stop a full data migration task, the data that is migrated to the RDS instance may be incomplete. You can wait until the full data migration task automatically stops.

• Incremental data migration

An incremental data migration task does not automatically stop. You must manually stop the task.

**?** Note We recommend that you manually stop an incremental data migration task at an appropriate point in time. For example, you can stop the task during off-peak hours or before you switch your workloads over to the RDS instance.

- a. Wait until Incremental Data Migration and The data migration task is not delayed appear in the progress bar of the data migration task. Then, stop writing data to the selfmanaged Oracle database for a few minutes. The delay time of incremental data migration may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The data migration task is not delayed** again. Then, manually stop the migration task.



#### FAQ

If an alert message appears when you set the RDS Instance ID parameter for the source database, you can solve the issue as described in following table.

Alert message	Solution		
Warning	Check whether the ID of the Alibaba Cloud account that owns the source instance and the role name configured for the account are valid.		
EntityNotExist.Role : The role not exists: acs:ram::1406926474064771;role/ram-for- dts1229. RequestId :	<b>Note</b> To obtain the ID of the Alibaba Cloud account that owns the source instance, you must log on to the Account Management console by using this account. The account ID is displayed on the Security Settings page.		
Warning X NoPermission : You are not authorized to do this action. You should be authorized by RAM. RequestId : CK	Make sure that you have performed the following steps: 1. Log on to the RAM console by using the Alibaba Cloud account that owns the source instance (Account A). 2. Specify the Alibaba Cloud account that owns the destination instance (Account B) as a trusted account. 3. Authorize Account B to access the cloud resources of Account A by using DTS. For more information, see Configure RAM authorization for cross-account data migration and synchronization.		

## 10.Migrate data from Alibaba Cloud to a self-managed database

### 10.1. Migrate data from an ApsaraDB RDS for MySQL instance to a selfmanaged MySQL database

This topic describes how to migrate data from an ApsaraDB RDS for MySQL instance to a self-managed MySQL database by using Data Transmission Service (DTS). This is applicable to scenarios such as data analysis and functional test.

#### Prerequisites

- The tables to be migrated from the ApsaraDB RDS for MySQL instance contain primary keys or UNIQUE NOT NULL indexes.
- The available storage space of the self-managed MySQL database is larger than the total size of the data in the ApsaraDB RDS for MySQL instance.
- The version of the self-managed MySQL database is the same as that of the ApsaraDB RDS for MySQL instance. This ensures compatibility.

#### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.
- If a data migration task fails, DTS automatically resumes the task. Therefore, before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination database after the task is resumed.

#### Limits

• DTS supports the following types of objects for schema migration: table, view, trigger, stored procedure, and function.

ONOTE During schema migration, DTS changes the value of the SECURITY attribute from DE FINER to INVOKER for views, stored procedures, and functions.

• DTS does not migrate user information from the source database. After data migration is complete, if you want to call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

#### **Migration types**

DTS supports schema migration, full data migration, and incremental data migration. For more information, see Terms.

**Note** When you migrate data between PolarDB for MySQL clusters, you can select all of the supported migration types to ensure service continuity.

#### Billing

Migration type	Task configuration fee	Internet traffic fee	
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.	

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

#### Permissions required for database accounts

Database	Required permissions
ApsaraDB RDS for MySQL	The read permission on the objects to be migrated
Self-managed MySQL database	The read and write permissions on the objects to be migrated

For more information about how to create and authorize a database account, see the following topics:

- ApsaraDB RDS for MySQL: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance
- Self-managed MySQL database: Create an account for a user-created MySQL database and configure binary logging

#### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

Region

1.Configure Source and Destination I	Databases 2.Configure Mi	igration Types and Objects	3.Map nam	e modification	4.Precheck
* Task Name:	RDS_TO_MySQL				
Source Database					
* Instance Type:	RDS Instance	•	DTS support type		
* Instance Region:	China (Hangzhou)	•			
* RDS Instance ID:	rm-b	•	RDS Instances of Other Apsara Stack Accounts		
* Database Account:	dtstest				
* Database Password:	•••••	4>	Test Connectivity	⊘ Passed	
* Encryption:	Non-encrypted O SSL-encrypte	d			
Destination Database					
* Instance Type:	User-Created Database in ECS In	stance v			
* Instance Region:	China (Hangzhou)	•			
* ECS Instance ID:	i-bp	•			
* Database Type:	MySQL	Ŧ			
* Port Number:	3306				
* Database Account:	dtstest				
* Database Password:	******	······ 《>		⊘ Passed	
* Encryption:	Non-encrypted OSL-encrypte	d			
					Cancel Set Whitelist and Next
Section	Parameter	Description			
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.			
	Instance Type	Select RDS Instance.			
	Instance	Select the region where the source RDS instance resides.			

Section	Parameter	Description		
	RDS Instance ID	Select the ID of the source RDS instance.		
	Dat abase Account	Enter the database account of the ApsaraDB RDS for MySQL instance. For information about the permissions that are required for the account, see Permissions required for database accounts.		
		Enter the password of the database account.		
Source Dat abase	Dat abase Password	<b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.		
	Encryption	Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you select <b>SSL-encrypted</b> , you must enable SSL encryption for the RDS instance before you configure the data migration task. For more information, see Configure SSL encryption for an ApsaraDB RDS for MySQL instance.		
		<b>Note</b> The <b>Encryption</b> parameter is available only for regions in mainland China and the China (Hong Kong) region.		
	Instance Type	Select an instance type based on the deployment of the destination database. In this example, select <b>User-Created Database in ECS Instance</b> .		
		<b>Note</b> If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see <b>Preparation overview</b> .		
	Instance Region	Select the region where the ECS instance resides.		
	ECS Instance ID	Select the ID of the ECS instance that hosts the self-managed MySQL database.		
	Database Type	Select MySQL.		
	Port Number	Enter the service port number of the self-managed MySQL database. In this example, enter <b>3306</b> .		
Destination				

Section	Parameter	Description	
	Dat abase Account	Enter the account of the self-managed MySQL database. For information about the permissions that are required for the account, see Permissions required for database accounts.	
	Dat abase Password	Enter the password of the database account.	
		<b>Note</b> After you specify the destination database parameters, click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Modify the destination database parameters based on the check results.	
	Encryption	Select Non-encrypted or SSL-encrypted. In this example, No encrypted is selected.	

#### 6. In the lower-right corner of the page, click Set Whitelist and Next.

Note DTS adds the CIDR blocks of DTS servers to the whitelist of the source ApsaraDB RDS for MySQL instance and the inbound rule of the destination ECS instance. This ensures that DTS servers can connect to the source and destination instances.

#### 7. Select the migration types and the objects to be migrated.

#### Dat a Transmission Service

1.Config	ure Source and	2.Configure Migration	Types and	3.Advanced Settings	>	4.Precheck
<ul> <li>Migration 1 triggers, For n</li> </ul>	Types: V Schema Migration nore information, see Refere	Full Data Migration	✓ Incremental I	Data Migration Note: Increm	ental data migrat	ion does not support
Note: do no cleans up ti	t clean up the incremental one log too early, the DTS inc	ata log generated by the sou remental task may fail	rce database after the	DTS task is started when the DT	'S full task is runn	ing. If the source database
Data migral between Ap For long-ter	tion applies to short-term mi osara Stack databases. m data synchronization in re	gration scenarios, Typical sce al time, use the data synchro	enarios include migratir onization feature.	ng data to the cloud, scaling and	sharding databasi	es, and migrating data
Available				Selected (To edit an object n Edit.) Learn more.	ame or its filter, ł	nover over the object and click
Expand the	e tree before you perform a	glol I <b>Q</b>				Q
🖂 📑 dtst	estdata			📼 dtstestdata (20bi	iects)	
E 📑	adies /iews			customer		
	It wiews		> <	i order		
Select All				Remove All		
*Rename Dat	abases and Tables:	Do Not Change Databa	se and Table Names	O Change Database and Ta	ble Names	
* Retry Time f	or Failed Connection	720 Minut	es 🕐			
*Source table DMS_ ONLINE_ Do you O Yes O No O the temporary table to the target database during DDL:						
Information: 1. Data migrat in the source o 2. Do not do D	ion only copies the data and database. DL operation during structu	schema in the source databa	ase and saves the copy se the task may fail	r in the destination database. The	e process does no	t affect any data or schema
				Cancel	Previous	Save Precheck
Setting	Description					
• To perform only full migration, select Schema Migration and Full Data Migrat				ata Migration.		
Select the	<ul> <li>To ensure s</li> <li>Data Migra</li> </ul>	ervice continuity ation, and Incre	during data i mental Data	migration, select Sc Migration.	hema Mi	gration, Full
migratio n types	Onte do not write consistency	If Incremental e data to the source between the source	Data Migrat Irce database urce and dest	t <b>ion</b> is not selected during data migra ination databases.	d, we recont tion. This (	mmend that you ensures data

#### Data Migration Migrate data from A libaba Cloud to a self-managed dat abase

Setting	Description		
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>		
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.		
Specify the retry time for failed connecti ons to the source or destinati on databas	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.		
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.		
е			

Setting	Description					
Specify whether to copy tempora ry tables to the destinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.					
on databas e when DMS perform s online DDL operatio ns on the source table	<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.					
	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.					
	<b>Note</b> If you select No, the tables in the destination database may be locked.					

#### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click **Next**.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.
Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



## 10.2. Migrate data from an ApsaraDB RDS for MySQL instance to a selfmanaged Kafka cluster

Kafka is a distributed message queue service that features high throughput and high scalability. Kafka is widely used for big data analytics such as log collection, monitoring data aggregation, streaming processing, and online and offline analysis. It is important for the big data ecosystem. This topic describes how to migrate data from an ApsaraDB RDS for MySQL instance to a self-managed Kafka cluster by using Data Transmission Service (DTS). The self-managed Kafka cluster has a public IP address. The data migration feature allows you to extend message processing capabilities.

### Prerequisites

- A Kafka cluster is created and the Kafka version is 0.10.1.0 to 2.7.0.
- The service port of the Kafka cluster must be accessible over the Internet.

### Context

If you use the data synchronization feature, the deployment of the self-managed Kafka cluster must belong to one of the following types:

- Self-managed database hosted on Elastic Compute Service (ECS)
- Self-managed database connected over Express Connect, VPN Gateway, or Smart Access Gateway
- Self-managed database without a public IP address or port number (connected over Database Gateway)
- Self-managed database connected over Cloud Enterprise Network (CEN)

If the deployment of your self-managed Kafka cluster does not belong to one of the preceding types, you can open the service port of the Kafka cluster to the Internet. Then, you can use the data migration feature to synchronize data to the Kafka cluster.

### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- You can select only tables as the objects to migrate.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

* Task Name: MySQL_TO_KAFKA							
Note: If you need to carry out incremental data migration for a long time, it is recommended to use the data synchronization function, which has better network stability and richer functions. Click to buy data synchronization Click to see the difference between data migration and synchronization							
Source Database							
* Instance Type:	RDS Instance 🗸	DTS support type					
* Instance Region:	China (Hangzhou)						
* RDS Instance ID:	rm-	RDS Instances of Other Apsara Stack Accounts					
* Database Account:	dtstest						
* Database Password:	Ø	Test Connectivity					
* Encryption:	● Non-encrypted ○ SSL-encrypted						
Destination Database							
* Instance Type:	User-Created Database with Public IP Address $\checkmark$						
* Instance Region:	China (Hangzhou)	Get IP Address Segment of DTS					
* Database Type:	Kafka 🗸						
* Hostname or IP Address:	172						
* Port Number:	9092						
Database Account:		Optional					
Database Password:	¢>	Optional					
* Topic:	dtstopic V Click Get Topic List and then select the specific topic.	Get Toplic list					
Topic for storing DDL:	~	Get Toplic list					
* Kafka Version	Click Get Topic List and then select the specific topic.						
* Encryption:	Non-encrypted O SCRAM-SHA-256						
* Whether to use Kafka schema registry:	● No ○ Yes						
		Cancel Set Whitelist and Next					

Section	Parameter	Description
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.
	lnstance Type	Select RDS Instance.
	Instance Region	The region where the source ApsaraDB RDS instance resides.
	Instance ID	The ID of the source ApsaraDB RDS instance.

Section	Parameter	Description		
Source	Database Account	The database account of the source instance. The account must have the SELECT permission on the objects to migrate and the REPLICATION CLIENT, REPLICATION SLAVE, and SHOW VIEW permissions.		
Database	Database Password	The password of the database account.		
	Encryption	Specifies whether to encrypt the connection to the source instance. Select <b>Non-encrypted</b> or <b>SSL-encrypted</b> . If you select <b>SSL-encrypted</b> , you must enable SSL encryption for the ApsaraDB RDS for MySQL instance before you configure the data migration task. For more information, see Configure SSL encryption for an ApsaraDB RDS for MySQL instance.		
		<b>Notice</b> The <b>Encryption</b> parameter is available only for regions in the Chinese mainland and the China (Hong Kong) region.		
	Instance Type	Select User-Created Database with Public IP Address.		
	Instance Region	You do not need to specify this parameter.		
	Database Type	Select Kafka.		
	Hostname or IP Address	The IP address that is used to access the self-managed Kafka cluster. In this example, the public IP address is used.		
	Port Number	The service port number of the Kafka cluster. Default value: 9092.		
Destination Database	Dat abase Account	The username that is used to log on to the Kafka cluster. If no authentication is enabled for the Kafka cluster, you do not need to enter the username.		
	Database Password	The password of the username. If no authentication is enabled for the Kafka cluster, you do not need to enter the password.		
	Торіс	Click <b>Get Topic List</b> , and select a topic name from the drop-down list.		
	Kafka version	The version of the destination Kafka cluster.		
	Encryption	Specifies whether to encrypt the connection to the destination cluster. Select <b>Non-encrypted</b> or <b>SCRAM-SHA-256</b> based on your business and security requirements.		

6. In the lower-right corner of the page, click **Set Whitelist and Next**.

#### □ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types, the migration policy, and the objects to migrate.

### Dat a Transmission Service

* Migration T triggers. For m	ypes: 🔽 Schema Mig ore information, see R	gration Reference	🗸 Full Data Mi	gration [	Incremental [	ata Migration	Note: Incre	mental data migr	ation does not s	support
Note: do no cleans up th Data migrat between Ap For long-ten	t clean up the increme ie log too early, the DT ion applies to short-te sara Stack databases, m data synchronization	ental data log TS increment rm migration n in real time	g generated by tal task may fail n scenarios. Typ e, use the data	the source da l iical scenarios synchronizatio	tabase after the include migratin on feature.	DTS task is star	ted when the [ oud, scaling an	DTS full task is rui d sharding datab:	nning. If the sou	urce database ting data
Available						Selected (To Edit.) Learn	edit an object	name or its filter	, hover over the	object and click
Expand the	estdata ables 'iews	rm a gloi			> <	iii cust iii cust iii orde	stdata (20) omer er	bjects)	Q	
Select All						Demous All				
*Rename Data	bases and Tables:	ا ھ	Do Not Change	Database and	l Table Names	Change [	Database and T	Table Names		
* Retry Time fi *Source table I want to copy th the target data	or Failed Connection DMS_ ONLINE_ Do yo he temporary table to ibase during DDL:	7 u 0 1	20 Yes 💿 No	Minutes (	D					
Information: 1. Data migrati in the source d 2. Do not do D	ion only copies the dat latabase. DL operation during st	ta and schen tructure and	na in the source full migration, (	e database an otherwise the	d saves the copy task may fail	in the destination	on database. T	he process does i	not affect any d	ata or schema
							Cancel	Previous	Save	Precheck

	Select Schema Migration, Full Data Migration, and Incremental Data Migration.
Select the migratio n types	<b>Notice</b> If <b>Incremental Data Migration</b> is not selected, we recommend that you do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases.
Select the data format used in Kafka	The data that is migrated to the Kafka cluster is stored in the Avro or Canal JSON format. For more information, see <mark>Data formats of a Kafka cluster</mark> .

Setting	Description
Select the policy for migratin g data to Kafka partition s	Select a migration policy based on your business requirements. For more information, see Specify the policy for synchronizing data to Kafka partitions.
Select the	Select one or more tables from the <b>Available</b> section and click the <b>&gt;</b> icon to add the tables to the <b>Selected</b> section.
objects that you want to migrate	<ul> <li>Note DTS maps the table names to the topic name that you select in Step 5.</li> <li>For information about how to rename the topic, see Object name mapping.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time range for failed connecti ons to the source or destinati on	By default, if DTS fails to connect to the source or destination database, DTS retries within the following 12 hours. You can specify the retry time range based on your business requirements. If DTS is reconnected to the source and destination databases within the specified time range, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> Within the time range in which DTS attempts to reconnect to the source and destination databases, you are charged for the DTS instance. We recommend that you specify the retry time range based on your business requirements. You can also release the DTS instance at the earliest opportunity after the source and destination databases are released.
e	

Setting	Description
Specify whether to copy tempora ry tables to the destinati	If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.
on databas e when DMS	<b>Note</b> If online DDL operations generate a large amount of data, latency may occur for the migration task.
perform s online	• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.
operatio ns on the source	<b>Note</b> If you select No, the tables in the destination database may be locked.
CANC	

### 8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



## 10.3. Migrate data from a PolarDB for MySQL cluster to a self-managed MySQL database

This topic describes how to migrate data from a PolarDB for MySQL cluster to a self-managed MySQL database by using Data Transmission Service (DTS). This is applicable to scenarios such as data analysis and functional test.

## Prerequisites

The binary logging feature is enabled for the source PolarDB for MySQL cluster. For more information, see Enable binlogging.

### Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- During full data migration, concurrent INSERT operations cause fragmentation in the tables of the destination database. After full data migration is complete, the size of used tablespace of the destination database is larger than that of the source database.
- If a data migration task fails, DTS automatically resumes the task. Therefore, before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database overwrites the data in the destination database after the task is resumed.

## Limits

• DTS supports the following types of objects for schema migration: table, view, trigger, stored procedure, and function.

**Note** During schema migration, DTS changes the value of the SECURITY attribute from DE FINER to INVOKER for views, stored procedures, and functions.

• DTS does not migrate user information from the source database. After data migration is complete, if you want to call a view, stored procedure, or function of the destination database, you must grant the read and write permissions to INVOKER.

## **Migration types**

DTS supports schema migration, full data migration, and incremental data migration. For more information, see Terms.

**Note** When you migrate data between PolarDB for MySQL clusters, you can select all of the supported migration types to ensure service continuity.

## Billing

Migration type	Task configuration fee	Internet traffic fee		
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from		
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.		

# SQL operations that can be synchronized during incremental data migration

Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul> <li>ALTER TABLE and ALTER VIEW</li> <li>CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW</li> <li>DROP INDEX and DROP TABLE</li> <li>RENAME TABLE</li> <li>TRUNCATE TABLE</li> </ul>

### Permissions required for database accounts

Database	Required permissions
PolarDB for MySQL	The read permissions on the objects to be migrated

Database	Required permissions
Self-managed MySQL database	The read and write permissions on the objects to be migrated

For more information about how to create and authorize a database account, see Create database accounts and Create an account for a self-managed MySQL database and configure binary logging.

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

1.Configure Source and Destination D	atabases 2.Configure Migration Types and Objects	$\rightarrow$	3.Map name	modification	$\rightarrow$	4.Precheck
* Task Name: P	olarDB_TO_MySQL	]				
Source Database						
* Instance Type:	PolarDB	۳	DTS support type			
* Instance Region:	China (Hangzhou)	۳				
* PolarDB Instance ID:	1. An artist state	•				
* Database Account:	dtstest					
* Database Password:	••••••		Test Connectivity	Passed		
Destination Database						
* Instance Type:	User-Created Database in ECS Instance	*				
* Instance Region:	China (Hangzhou)	•				
* ECS Instance ID:	i-bp	•				
* Database Type:	MySQL	•				
* Port Number:	3306					
* Database Account:	dtstest					
* Database Password:	*******	<⊅	Test Connectivity	Passed		
* Encryption:	Non-encrypted      SSL-encrypted					
						Cancel Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select PolarDB.
	Instance Region	Select the region where the source PolarDB cluster resides.

Section	Parameter	Description
	PolarDB Instance ID	Select the ID of the source PolarDB cluster.
	Dat abase Account	Enter the database account of the source PolarDB cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.
Source Database		Enter the password of the database account.
Dat abas Passwo	Dat abase Password	Note After you specify the source database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the source database parameters based on the check results.
		Select an instance type based on the deployment of the source database. In this example, select <b>User-Created Database in ECS Instance</b> .
	Instance Type	<b>Note</b> If you select other instance types, you must deploy the network environment for the source database. For more information, see Preparation overview.
	Instance Region	Select the region where the ECS instance resides.
	ECS Instance ID	Select the ID of the ECS instance that is connected to the self- managed MySQL database.
	Database Type	Select MySQL.
	Port Number	Enter the service port number of the self-managed MySQL database. In this example, enter <b>3306</b> .
Destination Database	Dat abase Account	Enter the account of the self-managed MySQL database. For information about the permissions that are required for the account, see Permissions required for database accounts.

Section	Parameter	Description
		Enter the password of the database account.
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the destination database parameters based on the check results.
	Encryption	Select Non-encrypted or SSL-encrypted. In this example, Non- encrypted is selected.

### 6. In the lower-right corner of the page, click Set Whitelist and Next.

**?** Note DTS adds the CIDR blocks of DTS servers to the whitelist of the source PolarDB for MySQL cluster and the inbound rule of the destination ECS instance. This ensures that DTS servers can connect to the source cluster and the destination instance.

### 7. Select the migration types and the objects to be migrated.

### Dat a Transmission Service

1.Config	ure Source and	2.Configure Migr	ation Types and	3.Advanced Settings	>	4.Precheck
* Migration T triggers. For m	Types: V Schema Migration nore information, see Refere	Full Data Mi	igration 🔽 Incremental	Data Migration Note: Incren	mental data migra	tion does not support
Note: do no cleans up th	t clean up the incremental d ne log too early, the DTS inc	ata log generated by remental task may fail	the source database after the I	e DTS task is started when the D	TS full task is run	ning. If the source database
Data migral between Ap For long-ter	tion applies to short-term mi Isara Stack databases. Im data synchronization in re	gration scenarios. Typ al time, use the data	vical scenarios include migrati synchronization feature.	ng data to the doud, scaling and	l sharding databas	ies, and migrating data
Available				Selected (To edit an object r Edit.) Learn more.	name or its filter,	hover over the object and click
Expand the	e tree before you perform a	glol Q				Q
🖂 🧰 dtst	estdata Tablee			📄 dtstestdata (20b	jects)	
	/iews			customer		
			> <	i order		
Select All				Remove All		
*Rename Data	abases and Tables:	Do Not Change	Database and Table Names	<ul> <li>Change Database and Ta</li> </ul>	able Names	
* Retry Time f	or Failed Connection	720	Minutes 🕜			
*Source table want to copy t the target data	DMS_ ONLINE_ Do you he temporary table to abase during DDL:	🔾 Yes ( No	0			
Information: 1. Data migrat in the source o 2. Do not do D	ion only copies the data and latabase. VDL operation during structu	schema in the source re and full migration, (	e database and saves the cop otherwise the task may fail	y in the destination database. Th	e process does no	st affect any data or schema
				Cancel	Previous	Save Precheck
Setting	Description					
	• To perform	only full mig	gration, select <b>Sc</b>	hema Migration a	and Full D	ata Migration.
Select the	<ul> <li>To ensure s</li> <li>Data Migra</li> </ul>	ervice contin ation, and In	uity during data cremental Data	migration, select So a Migration.	chema Mi	gration, Full
migratio n types	migratio n types Note If Incremental Data Migration is not selected, we recommend that you do not write data to the source database during data migration. This ensures data consistency between the source and destination databases.					

### Data Migration Migrate data from A libaba Cloud to a self-managed dat abase

Setting	Description
Select the objects to be migrate d	<ul> <li>Select one or more objects from the Available section and click the &gt; icon to move the objects to the Selected section.</li> <li>Note <ul> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination database, the name of the object remains unchanged. You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul> </li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination instance. For more information, see Object name mapping.
Specify the retry time for failed connecti ons to	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 720 minutes (12 hours). You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
the source or destinati on databas e	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.
ons to the source or destinati on databas e	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

Description			
If you use Data Management (DMS) to perform online DDL operations on the source database, you can specify whether to migrate temporary tables generated by online DDL operations. • Yes: DTS migrates the data of temporary tables generated by online DDL operations.			
<b>Note</b> If online DDL operations generate a large amount of data, the data migration task may be delayed.			
• <b>No</b> : DTS does not migrate the data of temporary tables generated by online DDL operations. Only the original DDL data of the source database is migrated.			
<b>Note</b> If you select No, the tables in the destination database may be locked.			

### 8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of **incremental data migration** changes to **The migration task is not delayed** again. Then, manually stop the migration task.



## 10.4. Migrate data from a PolarDB for PostgreSQL cluster to a self-managed Oracle database

This topic describes how to migrate data from a PolarDB for PostgreSQL cluster to a self-managed Oracle database by using Data Transmission Service (DTS). This is suitable for scenarios such as data reflow tests and functional tests.

## Prerequisites

- The tables to migrate from the source PolarDB for PostgreSQL cluster contain primary keys or UNIQUE NOT NULL indexes.
- The version number of the self-managed Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.
- The schemas of objects such as tables are created in the self-managed Oracle database.
- The available storage space of the self-managed Oracle database is larger than the total size of the data in the PolarDB for PostgreSQL cluster.

### Precautions

- In this scenario, DTS supports only full data migration and incremental data migration. DTS does not support schema migration.
- During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the loads of the database servers. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

- A data migration task can migrate data from only a single database. To migrate data from multiple databases, you must create a data migration task for each database.
- During incremental data migration, if you select a schema as the object to migrate, take note of the following limits: If you create a table in the schema or execute the RENAME statement to rename the table, you must execute the ALTER TABLE schema.table REPLICA IDENTITY FULL; Statement before you write data to the table.

Onte Replace the schema and table in the preceding sample statement with the actual schema name and table name.

• To ensure that the latency of incremental data migration is accurate, DTS adds a heartbeat table named <a href="https://dts\_postgres\_heartbeat">dts\_postgres\_heartbeat</a> to the source database. The following figure shows the schema of the heartbeat table.

=> select * slot_name	from dts_postgres_head   revice_time   revi	rtbeat; ice_lsn   flushed_	lsn   update_time	dts_service_time
v ai (1 row)	1617   194   0/4[	6 F0	2021-04-02 06:58:38.21	9416+00   161734 3

## Billing

Migration type	Task configuration fee	Internet traffic fee	
Full data migration	Free of charge.	Charged only when data is migrated from	
Incremental data migration	Charged. For more information, see Pricing.	Alibaba Cloud over the Internet. For more information, see Pricing.	

# SQL operations that can be synchronized during incremental data migration

INSERT, UPDATE, and DELETE

### Permissions required for database accounts

Database	Required permission
PolarDB for PostgreSQL cluster	Permissions of a privileged account
Self-managed Oracle database	Permissions of the schema owner

### Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click Create Migration Task.

5. Configure the source and destination databases.

* Tack Nama			
Task Natives	polarcopg_oracie ote: If you need to migrate incremental data for a long period of time	, we reci	ommend that you use the data synchronization feature, which has better
network stability. Click to buy data synchron	nization Click to see the difference between data migration and sync	hronizati	on
Source Database			
* Instance Tyre	ReleaDD		Supported Databases
***		•	
* Instance Region	China (Hangzhou)	~	
* PolarDB Instance ID	pr	•	in the other data bases for our first of shares. On the Minus Dame
* Data have been	ADAM evaluates the compatibility and transformation plan of Oracl	e migrao	on to other databases for you nee of charge do to view bemo
Database Name	dtstestdata		]
* Database Account	a dtstest		
* Database Password	·····	Ф	Test Connectivity
Destination Database			
* Instance Type	User-Greated Database with Public IP Address	~	]
* Instance Region	China (Hangzhou)	$\sim$	Get IP Address Segment of DTS
* Database Type	:: Oracle	~	
* Hostname or IP Address	52		]
* Port Number	1521		]
* Instance Type	• • Non-RAC Instance ORAC or PDB Instance		
* SID:			]
* Database Account			]
* Database Password	l:	Ф	Test Connectivity

Section	Parameter	Description
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.
	Instance Type	Select PolarDB.
	Instance Region	The region where the source PolarDB cluster resides.
	PolarDB Instance ID	The ID of the source PolarDB for PostgreSQL cluster.
	Dat abase Account	The database account of the source cluster. For information about the permissions that are required for the account, see Permissions required for database accounts.

Section Source	Parameter	Description
Database	Dat abase Password	The password of the database account. <b>Note</b> After you specify the source database parameters, click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to verify whether the specified parameters are valid. If the specified parameters are valid, the <b>Passed</b> message is displayed. If the <b>Failed</b> message is displayed, click <b>Check</b> next to <b>Failed</b> . Modify the source database parameters based on the check results.
	Instance Type	The access method of the destination self-managed database. In this example, <b>Public IP Address</b> is selected.           ⑦ Note       If the self-managed database is connected over other methods, you must set up the environment that is required for the database. For more information, see         Preparation overview.
	Instance Region	You do not need to specify this parameter.
	Database Type	Select <b>Oracle</b> .
	Hostname or IP Address	The IP address that is used to access the self-managed Oracle database. In this example, enter the public IP address.
Dectination	Port Number	The service port number of the self-managed Oracle database. In this example, <b>1521</b> is used.
Database	Instance Type	<ul> <li>If you select Non-RAC Instance, you must configure the SID parameter.</li> <li>If you select RAC or PDB Instance, you must specify the Service Name parameter.</li> <li>In this example, select Non-RAC Instance.</li> </ul>
	SID	The system ID (SID) of the destination database.

Section	Parameter	Description
	Dat abase Account	The account of the self-managed Oracle database. For information about the permissions that are required for the account, see Permissions required for database accounts.
		The password of the database account.
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the specified parameters are valid, the Passed message is displayed. If the Failed message is displayed, click Check next to Failed. Modify the destination database parameters based on the check results.

### 6. In the lower-right corner of the page, click Set Whitelist and Next.

### ☐ Warning

- If the source or destination database instance is an Alibaba Cloud database instance, such as an ApsaraDB RDS for MySQL or ApsaraDB for MongoDB instance, or is a self-managed database hosted on Elastic Compute Service (ECS), DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance or ECS security group rules. For more information, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases. If the source or destination database is a self-managed database on data centers or is from other cloud service providers, you must manually add the CIDR blocks of DTS servers to allow DTS to access the database.
- If the CIDR blocks of DTS servers are automatically or manually added to the whitelist of the database instance or ECS security group rules, security risks may arise. Therefore, before you use DTS to migrate data, you must understand and acknowledge the potential risks and take preventive measures, including but not limited to the following measures: enhance the security of your account and password, limit the ports that are exposed, authenticate API calls, regularly check the whitelist or ECS security group rules and forbid unauthorized CIDR blocks, or connect the database to DTS by using Express Connect, VPN Gateway, or Smart Access Gateway.
- After the DTS task is completed or released, we recommend that you manually detect and remove the added CIDR blocks from the whitelist of the database instance or ECS security group rules.
- 7. Select the migration types, the migration policy, and the objects to migrate.

abase

1.Configure So	urce and Destination 2.Configure Migration Types a	nd	3.Advanced Settings	$\rightarrow$	4.Precheck	
* Mi Refer	igration Types: V Full Data Migration V Incremental Data Migratience	on Note: Incre	emental data migration does not su	ipport triggers. For r	nore information, see	
No cle	te: do not clean up the incremental data log generated by the source d ans up the log too early, the DTS incremental task may fail	atabase after the D	/TS task is started when the DTS fu	ıll task is running. If	the source database	
Ava	ilable		Selected (To edit an object name Edit.) Learn more.	e or its filter, hover o	over the object and click	Hover over t
	xpand the tree before you perform a glo   Q dtstest Tables ♥ Views ♥ public	> <	dtstest (10bjects)	Q	]	the dialog be appears, moo object name destination d select the co migrate.
Sele *Ren: * Ret	ect All ame Databases and Tables: ry Time for Failed Connection 720 Minutes	d Table Names	Remove All O Change Database and Table	Names		
Infor 1. Dai in the 2. Do	mation: ta migration only copies the data and schema in the source database ar source database. not do DDL operation during structure and full migration, otherwise the	nd saves the copy i a task may fail	n the destination database. The pr	ccess does not affec	t any data or schema	echeck
						F
Paramet er	Description					
Rename Dat abas es and T ables	You can use the object name map the destination instance. For more	oing featu informati	re to rename the o on, see <mark>Object nan</mark>	objects tha ne mappin	at are migrate g.	ed to
Retry Time for	By default, if DTS fails to connect to the next 720 minutes (12 hours). Yo If DTS reconnects to the source an DTS resumes the data migration ta	to the sou ou can spe d destinat ask. Other	rce or destination ecify the retry time ion databases wit wise, the data mig	database, e range bas hin the spe ration task	DTS retries v sed on your n ecified time p c fails.	vithin leeds. leriod,
Failed Connecti on	<b>Note</b> When DTS retries a recommend that you specify the can also release the DTS instance destination instances are release	connectio e retry time e at your e ed.	n, you are charged e range based on y earliest opportunit	l for the Dī your busine y after the	S instance. V ess needs. Yc source and	Ve Du

8. In the lower-right corner of the page, click **Precheck**.

### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.
- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click Buy and Start to start the data migration task.
  - Full data migration

We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database may be incomplete. You can wait until the full data migration task automatically stops.

Full data migration and incremental data migration
 An incremental data migration task does not automatically stop. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** are displayed in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The latency of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the state of incremental data migration changes to The data migration task is not delayed again. Then, manually stop the migration task.

## 11.Migrate data between selfmanaged databases

## 11.1. Migrate data between selfmanaged Oracle databases

This topic describes how to migrate data between self-managed Oracle databases by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. You can select all of the supported migration types to ensure service continuity.

## Prerequisites

• The versions of the source and destination Oracle databases are 9i, 10g, 11g, 12c, 18c, or 19c.

(?) Note To ensure compatibility, make sure that the versions of the source and destination databases are the same.

- Supplement al logging, including SUPPLEMENT AL\_LOG\_DAT A\_PK and SUPPLEMENT AL\_LOG\_DAT A\_UI, is enabled for the source Oracle dat abase. For more information, see Supplement al Logging.
- The source Oracle database is running in ARCHIVELOG mode. Archived log files are accessible and a suitable retention period is set for archived log files. For more information, see Managing Archived Redo Log Files.
- The available storage space of the destination Oracle database is larger than the total size of the data in the source Oracle database.

## Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is unfavorable, the specification is low, or the data volume is large, database services may become unavailable. For example, DTS occupies a large amount of read and write resources in the following cases: a large number of slow SQL queries are performed on the source database, the tables have no primary keys, or a deadlock occurs in the destination database. Before you migrate data, evaluate the impact of data migration on the performance of the source and destination databases. We recommend that you migrate data during off-peak hours. For example, you can migrate data when the CPU utilization of the source and destination databases is less than 30%.
- The tables to be migrated in the source database must have PRIMARY KEY or UNIQUE constraints and all fields must be unique. Otherwise, the destination database may contain duplicate data records.
- If the self-managed Oracle database is deployed in a Real Application Cluster (RAC) architecture and is connected to DTS over an Alibaba Cloud virtual private cloud (VPC), you must connect the Single Client Access Name (SCAN) IP address of the Oracle RAC and the virtual IP address (VIP) of each node to the VPC and configure routes. The settings ensure that your DTS task can run as expected. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.

Notice When you configure the source Oracle database in the DTS console, you can specify the SCAN IP address of the Oracle RAC as the database endpoint or IP address.

• If a data migration task fails, DTS automatically resumes the task. Before you switch your workloads to the destination database, stop or release the data migration task. Otherwise, the data in the source database will overwrite the data in the destination database after the task is resumed.

## Billing

Migration type	Task configuration fee	Internet traffic fee
Schema migration and full data migration	Free of charge. Charged only when data is migrated f	
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

## Migration types

Migration type	Description
Schema migration	DTS migrates the schemas of the required objects to the destination Oracle database. DTS supports schema migration for the following types of objects: table, view, synonym, trigger, stored procedure, function, package, and user-defined type.
	<b>Note</b> If an object contains triggers, the data between the source and destination databases will become inconsistent.
Full data migration	DTS migrates historical data of the required objects from the source Oracle database to the destination Oracle database.
	<b>Note</b> During schema migration and full data migration, we recommend that you do not perform data definition language (DDL) operations on the required objects. Otherwise, the objects may fail to be migrated.
Incremental data migration	After full data migration, DTS retrieves redo log files from the source Oracle database. Then, DTS synchronizes incremental data from the source Oracle database to the destination Oracle database. Incremental data migration allows you to ensure service continuity when you migrate data between Oracle databases.

## Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Source Oracle database	The permissions of the schema owner	The permissions of the schema owner	The database administrator (DBA) permission
Destination Oracle database	The permissions of the schema owner	The permissions of the schema owner	The permissions of the schema owner

For more information about how to create and authorize an Oracle database account, see CREATE USER and GRANT.

### Procedure

The procedure in this topic uses a **self-managed database hosted on ECS** as an example. You can also follow the procedure to configure data migration tasks for other types of self-managed Oracle databases.

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination cluster resides.
- 4. In the upper-right corner of the page, click **Create Migration Task**.
- 5. Configure the source and destination databases.

* Tas	k Name:	Oracle_to_Oracle					
Source Database							
* Instan	ice Type:	User-Created Database in E	CS Instance 🔻	DTS support type			
* Instance	e Region:	China (Hangzhou)	Ŧ				
* ECS Inst	tance ID:	i-bp	•				
* Databa	se Type:	Oracle	٣				
* Port	Number:	1521					
* Instan	ice Type:	Non-RAC Instance ORAC	C Instance				
	* SID:	testsid		]			
* Database	Account:	dtstest					
* Database P	assword:	•••••	4>	Test Connectivity	⊘ Passed		
Destination Database							
* Instan	ice Type:	User-Created Database in E	CS Instance 🔹				
* Instance	e Region:	China (Hangzhou)	•				
* ECS Inst	ance ID:	i-bp	•				
* Databa	* Database Type: Oracle •						
* Port	Number:	1521					
* Instan	ice Type:	Non-RAC Instance ORAC	Instance				
	* SID:	testsid1					
* Database	Account:	dtstest					
* Database P	* Database Password:		4>	Test Connectivity	⊘ Passed		
					Cancel	Assess Data Migration to Cloud	Set Whitelist and Next
Section	Par	ameter	Description				

N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name to identify the task. You do not need to use a unique task name.

Section	Parameter	Description
	Instance Type	Select User-Created Database in ECS Instance.  Note If you select other instance types, you must deploy the network environment for the source database. For more information, see Preparation overview.
	Instance Region	Select the region of the Elastic Compute Service (ECS) instance on which the source Oracle database is deployed.
	ECS Instance ID	Select the ID of the ECS instance on which the source Oracle database is deployed.
	Database Type	Select Oracle.
	Port Number	Enter the service port number of the source Oracle database.
Source Database Ins Da Ac	Instance Type	<ul> <li>If you select Non-RAC Instance, you must specify the SID parameter.</li> <li>If you select RAC Instance, you must specify the Service Name parameter.</li> </ul>
	Dat abase Account	Enter the account of the source Oracle database. For information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the database account.
	Dat abase Password	<b>?</b> Note After you specify the information about the self- managed Oracle database, you can click <b>Test Connectivity</b> next to <b>Database Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.
	Instance Type	Select User-Created Database in ECS Instance.
	Instance Region	Select the region of the ECS instance on which the destination Oracle database is deployed.
	ECS Instance ID	Select the ID of the ECS instance on which the destination Oracle database is deployed.
	Database Type	Select Oracle.
	Port Number	Enter the service port number of the destination Oracle database.

Section	Parameter	Description		
Destinatio n Database	Instance Type	<ul> <li>If you select Non-RAC Instance, you must specify the SID parameter.</li> <li>If you select RAC Instance, you must specify the Service Name parameter.</li> </ul>		
	Dat abase Account	Enter the account of the destination Oracle database. For information about the permissions that are required for the account, see <b>Permissions required for database accounts</b> .		
	Dat abase Password	Enter the password of the database account.		
		<b>Note</b> After you specify the information about the RDS instance, you can click <b>Test Connectivity</b> next to <b>Database</b> <b>Password</b> to check whether the information is valid. If the information is valid, the <b>Passed</b> message appears. If the <b>Failed</b> message appears, click <b>Check</b> next to <b>Failed</b> . Then, modify the information based on the check results.		

- 6. In the lower-right corner of the page, click Set Whitelist and Next.
- 7. Select the migration types and the objects to be migrated.

### Data Migration Migrate data betwe

### en self-managed dat abases

Available			
			Selected (To edit an object name or its filter, hover over the object and o Edit.) Learn more.
Expand the tree before you perform a	a glol 🛛 🔍 🔍		
🖃 🧰 test123			📔 dtstestdata
		>	
		<	
Solort All			Remove All
Jaccini			Renove All
*Rename Databases and Tables:	Do Not Change Database an	d Table Names	Change Database and Table Names
*Rename Databases and Tables: * Retry Time for Failed Connection	<ul> <li>Do Not Change Database an</li> <li>720 Minutes</li> </ul>	nd Table Names	Change Database and Table Names

Setting	Description
Solart the	<ul> <li>To perform only full data migration, select Schema Migration and Full Data Migration.</li> <li>To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration.</li> </ul>
Select the migration types	<b>Note</b> If <b>Incremental Data Migration</b> is not selected, do not write data to the source Oracle database during full data migration. This ensures data consistency between the source and destination databases. For more information, see <b>Migration types</b> .

Setting	Description
Select the objects to be migrated	Select one or more objects from the <b>Available</b> section and click the > icon to move the objects to the <b>Selected</b> section.
	<ul> <li>Note</li> <li>You can select columns, tables, or databases as the objects to be migrated.</li> <li>By default, after an object is migrated to the destination Oracle database, the name of the object remains the same as that in the source Oracle database. You can use the object name mapping feature to rename the objects that are migrated to the destination Oracle database. For more information, see Object name mapping.</li> <li>If you use the object name mapping feature to rename an object, other objects that are dependent on the object may fail to be migrated.</li> </ul>
Specify whether to rename objects	You can use the object name mapping feature to rename the objects that are migrated to the destination database. For more information, see Object name mapping.
Specify the retry time for failed connections to the source or destination database	By default, if DTS fails to connect to the source or destination database, DTS retries within the next 12 hours. You can specify the retry time based on your needs. If DTS reconnects to the source and destination databases within the specified time, DTS resumes the data migration task. Otherwise, the data migration task fails.
	<b>Note</b> When DTS retries a connection, you are charged for the DTS instance. We recommend that you specify the retry time based on your business needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

### 8. In the lower-right corner of the page, click **Precheck**.

#### ? Note

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- If the task fails to pass the precheck, you can click the *icon* icon next to each failed item

to view details.

- You can trouble shoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click Next.

- 10. In the **Confirm Settings** dialog box, specify the **Channel Specification** parameter and select **Data Transmission Service (Pay-As-You-Go) Service Terms**.
- 11. Click **Buy and Start** to start the data migration task.
  - Schema migration and full data migration
     We recommend that you do not manually stop the task during full data migration. Otherwise, the data migrated to the destination database will be incomplete. You can wait until the data migration task automatically stops.
  - Schema migration, full data migration, and incremental data migration
     The task does not automatically stop during incremental data migration. You must manually stop the task.

Notice We recommend that you select an appropriate time to manually stop the data migration task. For example, you can stop the task during off-peak hours or before you switch your workloads to the destination cluster.

- a. Wait until **Incremental Data Migration** and **The migration task is not delayed** appear in the progress bar of the migration task. Then, stop writing data to the source database for a few minutes. The delay time of **incremental data migration** may be displayed in the progress bar.
- b. Wait until the status of incremental data migration changes to The migration task is not delayed again. Then, manually stop the migration task.

