Alibaba Cloud

Data Transmission Service

Data Migration

Document Version: 20220617

(-) 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 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> Danger</u>	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.	
<u> </u>	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.	
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 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 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 {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}	

Table of Contents

1.Data migration procedure list	08
2.ETL features	11
2.1. Data type conversion	11
2.2. Object name mapping	15
2.3. Data filtering	21
3.Throttling	25
3.1. Configure throttling for full data migration	25
4.Migration task management	26
4.1. View the progress of a data migration task	26
4.2. View performance metrics of full data migration	28
4.3. View performance metrics of incremental data migration	29
4.4. Restart a failed data migration task	30
4.5. Stop a data migration task	32
5.Application switchover	34
5.1. Switch workloads to the target database	34
6.Precheck troubleshooting	37
6.1. Source database connectivity	37
6.2. Source database binary logging	38
6.3. Integrity constraints	41
6.4. MySQL storage engine	43
6.5. Source database permissions	43
6.6. Target database permissions	43
6.7. Object name conflicts	44
6.8. Target database availability	45
6.9. server_id in the source database	46
6.10. Target database connectivity	47

6.11. Source database engine version	48
7.Migrate from a user-created database to Alibaba Cloud	19
7.1. Source: MySQL	19
7.1.1. Migrate data from a user-created MySQL database to a	19
7.1.2. Migrate data from a user-created MySQL database conn	56
7.1.3. Migrate data from a user-created MySQL database conn	51
7.1.4. Migrate data from a user-created MySQL database to a	71
7.1.5. Migrate data from a user-created MySQL database to a	79
7.2. Source: SQL Server	36
7.2.1. Migrate incremental data from a user-created SQL Serv	36
7.2.2. Migrate full data from a user-created SQL Server datab	94
7.3. Source: Oracle 10)0
7.3.1. Migrate data from a user-created Oracle database to a 10)0
7.3.2. Migrate data from a user-created Oracle database to a 10)7
7.3.3. Migrate data from a user-created Oracle database to a 11	15
7.4. Source: PostgreSQL 12	22
7.4.1. Migrate incremental data from a user-created PostgreSQ 12	22
7.4.2. Migrate incremental data from a user-created PostgreS 13	30
7.4.3. Migrate full data from a user-created PostgreSQL data	38
7.5. Source: Redis 14	14
7.5.1. Migrate data from a user-created Redis database to an 14	14
7.6. Source: MongoDB	51
7.6.1. Migrate user-created standalone MongoDB databases to	51
7.6.2. Migrate the replica set of a user-created MongoDB dat	59
7.6.3. Migrate the user-created sharded MongoDB database t 16	56
7.7. Source: TiDB	76
7.7.1. Migrate incremental data from a user-created TiDB data 17	76
7.7.2. Migrate full data from a user-created TiDB database to 18	34

7	7.8. Source: Db2	191
	7.8.1. Migrate data from a user-created Db2 database to an	191
8.1	Migrate from a third-party cloud to Alibaba Cloud	199
8	3.1. Migrate data from an Amazon RDS for MySQL instance to	199
8	3.2. Migrate data from an Amazon RDS for Oracle database to	208
8	3.3. Migrate incremental data from an Amazon RDS for Postgr	218
8	3.4. Migrate full data from an Amazon RDS for PostgreSQL da	226
8	3.5. Migrate data from an Amazon Aurora MySQL database to	234
8	3.6. Migrate data from an Amazon Aurora MySQL database to	242
8	3.7. Migrate full data from an Amazon Aurora PostgreSQL data	250
8	3.8. Migrate full data from an Amazon RDS for SQL Server da	259
8	3.9. Migrate data from a MongoDB Atlas database to Alibaba	267
9.1	Migrate using the same Alibaba Cloud account	275
ç	9.1. Migrate data between RDS instances	275
ç	9.2. Migrate data from an ApsaraDB RDS for MariaDB TX insta	281
ç	9.3. Migrate data from an ApsaraDB RDS for MySQL instance	289
ç	9.4. Migrate data between Apsara PolarDB for MySQL clusters	296
9	9.5. Migrate data from an Apsara PolarDB for MySQL cluster t	302
ç	9.6. Migrate data between PolarDB clusters compatible with Or	309
ç	9.7. Migrate data from a standalone MongoDB instance to a re	315
ç	9.8. Migrate data from a replica set MongoDB instance to a sh	320
9	9.9. Migrate the data of an ApsaraDB for MongoDB instance a	326
10.	.Migrate across different Alibaba Cloud accounts	338
1	10.1. Migrate data between Apsara PolarDB for MySQL clusters	338
1	10.2. Migrate data between ApsaraDB for MongoDB instances	347
11.	Migrate from Alibaba Cloud to a user-created database	358
1	11.1. Migrate data from an ApsaraDB RDS for MySQL database	358
1	11.2. Migrate data from an Apsara PolarDB for MySQL cluster	362

12.Migrate between user-created databases	-	37
12.1. Migrate data between user-created Oracle databases	-	37

1.Data migration procedure list

The data migration mode migrates data between data stores. This mode is typically used for one-time migrations that require minimized downtimes. The data migration mode also provides transform, and load (ETL) features such as object name mapping and data filtering. The data migration mode allows you to migrate either into or out of Alibaba Cloud.

Note For more information about the supported sources and targets, see Supported sources and targets

Scenario	Source database	Topics
	MySQL	Migrate data from a user-created MySQL database to an ApsaraDB RDS for MySQL instance
		Migrate data from a user-created MySQL database connected over Express Connect, VPN Gateway, or Smart Access Gateway to an ApsaraDB RDS for MySQL database
		Migrate data from a user-created MySQL database connected over Express Connect, VPN Gateway, or Smart Access Gateway to an ApsaraDB RDS for MySQL instance across Alibaba Cloud accounts
		Migrate data from a user-created MySQL database to an Apsara PolarDB for MySQL cluster
Migrate data from a user- created database to Alibaba Cloud		Migrate data from a user-created MySQL database to a DRDS instance
	SQL Server	Migrate incremental data from a user-created SQL Server database to an ApsaraDB RDS for SQL Server instance
		Migrate full data from a user-created SQL Server database to ApsaraDB RDS for SQL Server
	Oracle	Migrate data from a user-created Oracle database to an ApsaraDB RDS for MySQL instance
		Migrate data from a user-created Oracle database to a DRDS instance
		Migrate data from a user-created Oracle database to an ApsaraDB RDS for PPAS instance
		Migrate data from a user-created Oracle database to a PolarDB cluster compatible with Oracle
		Migrate incremental data from a user-created PostgreSQL database (version 10.1 to 12) to an ApsaraDB RDS for PostgreSQL instance

Scenario	Source database	Topics
	PostgreSQL	Migrate incremental data from a user-created PostgreSQL database (version 9.4 to 10.0) to an ApsaraDB RDS for PostgreSQL instance
		Migrate full data from a user-created PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance
	Redis	Migrate data from a user-created Redis database to ar ApsaraDB for Redis instance
		Migrate user-created standalone MongoDB databases to Alibaba Cloud by using DTS
	MongoDB	Migrate the replica set of a user-created MongoDB database to ApsaraDB for MongoDB by using DTS
		Migrate the user-created sharded MongoDB database to ApsaraDB for MongoDB by using DTS
	TiDB	Migrate incremental data from a user-created TiDB database to an ApsaraDB RDS for MySQL instance
		Migrate full data from a user-created TiDB database to an ApsaraDB RDS for MySQL instance
	Db2	Migrate data from a user-created Db2 database to an ApsaraDB RDS for MySQL instance
	Amazon RDS	Migrate data from an Amazon RDS for MySQL instance to an ApsaraDB RDS for MySQL instance
		Migrate data from an Amazon RDS for Oracle database to an ApsaraDB RDS for MySQL database
		Migrate full data from an Amazon RDS for Oracle database to an ApsaraDB RDS for PPAS database
		Migrate full data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL databas
Migrate data from a third- party cloud to Alibaba	Amazon Aurora	Migrate data from an Amazon Aurora MySQL database to an ApsaraDB RDS for MySQL database
Cloud		Migrate data from an Amazon Aurora MySQL database to a PolarDB for MySQL database
		Migrate full data from an Amazon Aurora PostgreSQL database to an ApsaraDB RDS for PostgreSQL databas
	Amazon RDS for SQL Server	Migrate full data from an Amazon RDS for SQL Server database to an ApsaraDB RDS for SQL Server database

Scenario	Source database	Topics
	RDS instance	Migrate data between RDS instances
	ApsaraDB RDS for MySQL instance	Migrate data from an ApsaraDB RDS for MySQL instance to an Apsara PolarDB for MySQL cluster
	ApsaraDB RDS for MariaDB TX instance	Migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance
Migrate data between the same Alibaba Cloud	Apsara PolarDB for	Migrate data between Apsara PolarDB for MySQL clusters
account	MySQL cluster	Migrate data from an Apsara PolarDB for MySQL cluster to an ApsaraDB RDS for MySQL instance
	ApsaraDB for MongoDB instance	Migrate data from a standalone MongoDB instance to a replica set or sharded cluster instance
		Migrate data from a replica set MongoDB instance to a sharded cluster instance
		Migrate the data of an ApsaraDB for MongoDB instance across regions
	ApsaraDB for RDS instance	Migrate data between RDS instances across Alibaba Cloud accounts
Migrate data across Alibaba Cloud accounts	ApsaraDB for PolarDB cluster	Migrate data between Apsara PolarDB for MySQL clusters across Alibaba Cloud accounts
	ApsaraDB for MongoDB instance	Migrate data between ApsaraDB for MongoDB instances created by different Alibaba Cloud accounts
	ApsaraDB RDS for MySQL instance	Migrate data from an ApsaraDB RDS for MySQL database to a user-created MySQL database
Migrate data from Alibaba Cloud to a user-created database	ApsaraDB RDS for SQL Server instance	Migrate incremental data from an ApsaraDB RDS for SQL Server database to a user-created SQL Server database
	Apsara PolarDB for MySQL	Migrate data from an Apsara PolarDB for MySQL cluster to a user-created MySQL database
Migrate data between user-created databases	Oracle	Migrate data between user-created Oracle databases

2.ETL features

2.1. Data type conversion

For a heterogeneous migration, Data Transmission Service converts columns to data types supported by the target database engine. The following tables list the mappings of source and target types for supported migration paths.

TiDB-to-MySQL

Source: a user-created TiDB database

Target: a user-created MySQL database, an ApsaraDB RDS for MySQL database, or an ApsaraDB for PolarDB database

TiDB data type	MySQL data type
BIGINT	BIGINT
BINARY	BINARY
BIT	BIT
BOOL\ BOOLEAN	TINYINT
CHAR	CHAR
DATE	DATE
DATETIME	DATETIME
DECIMAL	DECIMAL
DOUBLE	DOUBLE
ENUM	ENUM
FLOAT	FLOAT
INT	INT
INTEGER	INT EGER
JSON	JSON
MEDIUMBLOB/LONGBLOB TINYBLOB/BLOB/	MEDIUMBLOB/LONGBLOB TINYBLOB/BLOB/
MEDIUMINT	MEDIUMINT
SET	SET
SMALLINT	SMALLINT

TiDB data type	MySQL data type
T EXT / LONGT EXT	TEXT/LONGTEXT
TIME	TIME
TIMESTAMP	TIMESTAMP
TINYINT	TINYINT
VARBINARY	VARBINARY
VARCHAR	VARCHAR
YEAR	YEAR

Oracle-to-MySQL

Source: a user-created Oracle database

Target: a user-created MySQL database, an ApsaraDB RDS for MySQL database, or an ApsaraDB for PolarDB database

Oracle data type	MySQL data type
varchar2(n [char/byte])	varchar(n)
nvarchar2[(n)]	national varchar[(n)]
char[(n [byte/char])]	char[(n)]
nchar[(n)]	national char[(n)]
number[(p[,s])]	decimal[(p[,s])]
float(p)]	double
long	longtext
date	datetime
binary_float	decimal(65,8)
binary_double	double
timestamp[(fractional_seconds_precision)]	datetime[(fractional_seconds_precision)]
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]
clob	longtext

Oracle data type	MySQL data type
nclob	longtext
blob	longblob
raw	varbinary(2000)
long raw	longblob
bfile	Not supported
interval year(year_precision) to month	Not supported
interval day(day_precision)to second[(fractional_seconds_precision)]	Not supported

? Note

- A CHAR column with a length greater than 255 bytes is converted to the varchar(n) type.
- Data types such as BFILE, INTERVAL YEAR TO MONTH, and INTERVAL DAY TO SECOND in Oracle databases are not supported in MySQL databases. They cannot be converted to data types supported by the destination database during schema migration.
 The schema migration fails if the table to be migrated contains unsupported data types.
 You must make sure that these unsupported columns are excluded from the objects to be migrated.
- The TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE data types in
 Oracle databases contain time zone information. However, the timestamp data type of
 MySQL databases does not contain time zone information. Therefore, DTS translates the
 time values of these data types into UTC time before transferring them to the target
 database.

Oracle-to-DRDS

Source: a user-created Oracle database

Target: a database created on a DRDS instance

Oracle data type	DRDS data type
varchar2(n [char/byte])	varchar(n)
nvarchar2[(n)]	national varchar[(n)]
char[(n [byte/char])]	char[(n)]
nchar[(n)]	national char[(n)]
number[(p[,s])]	decimal[(p[,s])]
float(p)]	double

Oracle data type	DRDS data type
oracle data type	
long	longtext
date	datetime
binary_float	decimal(65,8)
binary_double	double
timestamp[(fractional_seconds_precision)]	datetime[(fractional_seconds_precision)]
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]
clob	longtext
nclob	longtext
blob	longblob
raw	varbinary(2000)
long raw	longblob
bfile	Not supported
interval year(year_precision) to month	Not supported
<pre>interval day(day_precision)to second[(fractional_seconds_precision)]</pre>	Not supported

? Note

- A CHAR column with a length greater than 255 bytes is converted to the VARCHAR(n) type in the DRDS database.
- The TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE types in Oracle databases contain time zone information. However, the timestamp data type of DRDS databases does not contain the time zone information. Therefore, DTS translates the time values of these data types into UTC time before transferring them to the target database.

Oracle-to-PPAS

Source: a user-created Oracle database

Target: an ApsaraDB RDS for PPAS database

Oracle data type	PPAS data type
varchar2(n [char/byte])	varchar2[(n)]

Oracle data type	PPAS data type
nvarchar2[(n)]	nvarchar2[(n)]
char[(n [byte/char])]	char[(n)]
nchar[(n)]	nchar[(n)]
number[(p[,s])]	number[(p[,s])]
float(p)]	double precision
long	long
date	date
binary_float	real
binary_double	double precision
timestamp[(fractional_seconds_precision)]	timestamp[(fractional_seconds_precision)]
timestamp[(fractional_seconds_precision)]with time zone	timestamp[(fractional_seconds_precision)]with time zone
timestamp[(fractional_seconds_precision)]with local time zone	timestamp[(fractional_seconds_precision)]with time zone
clob	clob
nclob	nclob
blob	blob
raw	raw(size)
long raw	long raw
bfile	Not supported
interval year(year_precision) to month	interval year to month
interval day(day_precision) to second[(fractional_seconds_precision)]	interval day to second[(fractional_seconds_precision)]

Note The TIMESTAMP [(fractional_seconds_precision)] WITH LOCAL TIME ZONE data type in Oracle databases is not supported in ApsaraDB RDS for PPAS databases. DTS translates the time values of these data types into UTC time before storing them as the TIMESTAMP [(fractional_seconds_precision)] WITH TIME ZONE data type of the target instance.

2.2. Object name mapping

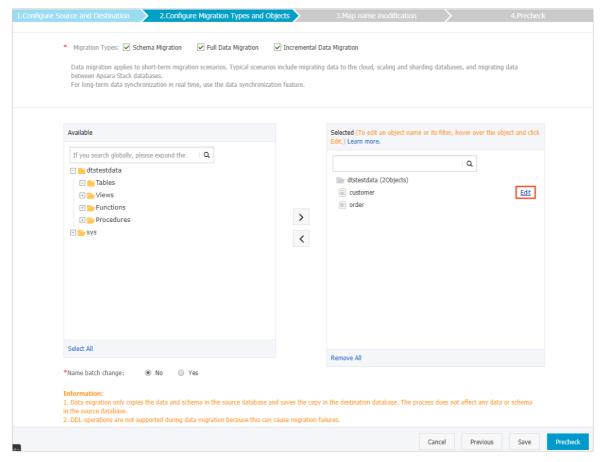
Data Transmission Service (DTS) allows you to configure name mappings between source and target columns, tables, or databases that have different names. You can either define one-to-one mappings for individual objects or specify renaming rules to rename multiple objects at a time.

Limits

MongoDB supports only database name mapping and table name mapping.

Configure name mappings for individual databases, tables, and columns

- 1. Create a data migration task.
- 2. In the **Configure Migration Types and Objects** step, select the required objects into the **Selected** list, move the pointer over a database or table, and then click **Edit** next to the object.

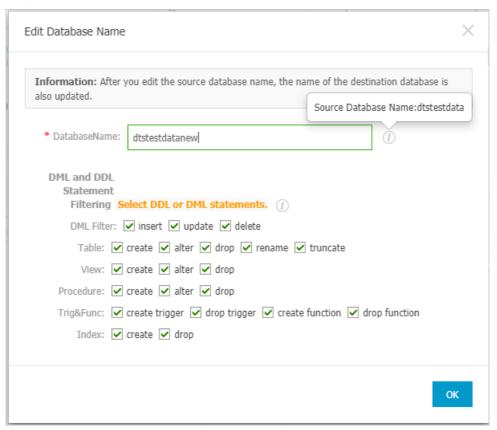


3. In the dialog box that appears, specify a name for the object in the target instance.

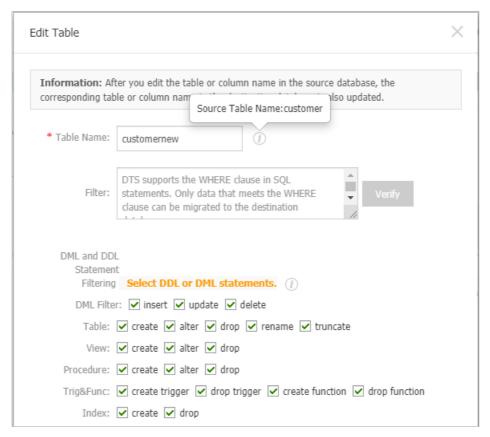
? Note

- If the source or target database is a user-created MySQL database, ApsaraDB RDS for MySQL, or Apsara PolarDB for MySQL, you can specify the data definition language (DDL) and data manipulation language (DML) operations that you want to replicate to the target database.
- If different sets of operations are selected for database name mapping and table name mapping, the operations that are selected for table name mapping take precedence.
- Database name mapping

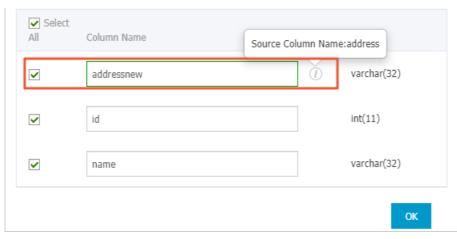
In the **Edit Database Name** dialog box, enter the database name that you want to use in the target database system.



• Table name mapping
In the **Edit Table** dialog box, enter the table name that you want to use in the target database.



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



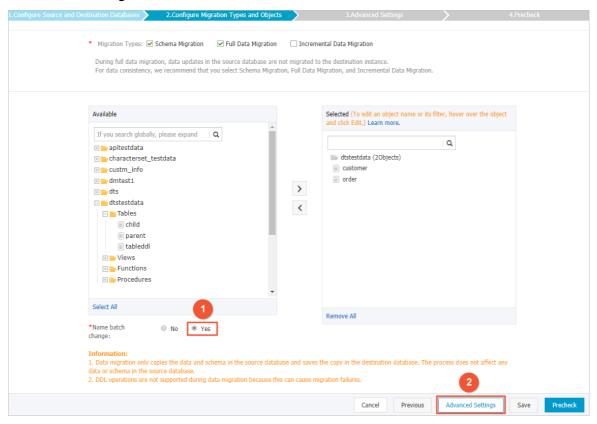
Note In this step, you can clear the selection for columns that you do not want to migrate.

- 4. Click OK.
- 5. Complete other configurations and continue to create a task.

Configure name mappings for multiple tables and columns at a time

- 1. Create a data migration task.
- 2. In the Configure Migration Types and Objects step, select the required objects into the Selected list.

- Note If the last object in your selection is a database, you cannot configure name mapping for tables or columns by using the Name batch change setting.
- 3. In the lower part of the page, set the **Name batch change** setting to **Yes**, and then click **Advanced Settings**.



- 4. In the Advanced Settings step, create renaming rules for tables and columns as follows:
 - i. For the **Choice Range** setting, select the scope of objects to which you want to apply the renaming rules. To select all objects, select **All**. To select specific objects, select **Section**, and then enter keywords to search for the desired databases (the text box in the left) and tables (the text box in the right).
 - ii. For the **Selection Rules** setting, select a renaming pattern from the drop-down list, and enter the corresponding naming element in the text box. The available renaming patterns include commonly used patterns, such as prefixing and suffixing.
 - iii. Click Add Modification Rules.
 - Note You can repeat steps i and ii to add more rules for different tables.

In this task, a total of 1 libraries and 2 tables are selected, and multiple rule conditions can be edited for name modification: Choice

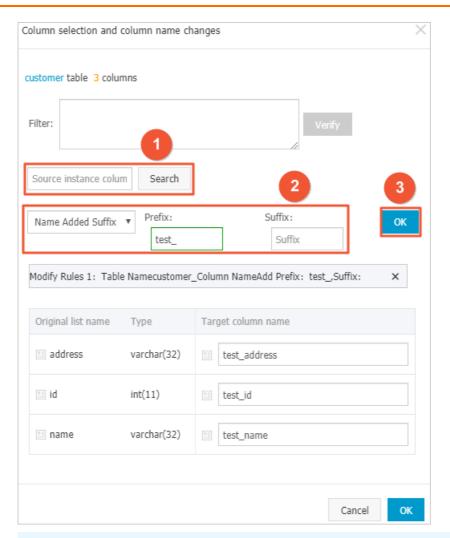
Section

All(All library entries) Range: Source instance library Source instance table A total of 1 libraries 2 tables that meet the above conditions Selection Change Name ▼ Change Name: Name Preview Modify Rules1: Table Name(customer)Add Prefix:test_,Suffix: X Modify Rules2: Table Name(order)Change Name: neworder X Overall Preview Partial Preview Source Database Destination Database(Click on the name below to edit directly) dtstestdata (2) Batch Change Table Name dtstestdata test_customer Column Name Modification customer Column Name Modification Save Precheck

iv. Click Preview to verify that the objects are renamed to the target names as expected.

- v. On the Partial Preview tab, click the icon next to the database name in the Destination

 Database section, and then click Column Name Modification.
- vi. Configure renaming rules for column names, and then click OK.



? Note

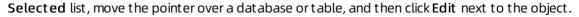
- By default, a renaming rule applies to all column names. You can search for a specific column name and modify the renaming rule. To do this, enter the column name in the search box, and then click Search.
- In the dialog box, you can specify conditions to filter data.
- vii. Click OK.
- 5. Click Precheck.
- 6. Complete other configurations and continue to create a task.

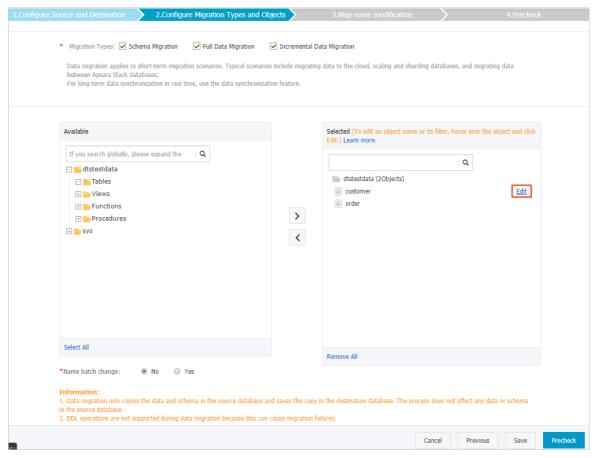
2.3. Data filtering

You can use SQL-like queries to filter the source table to migrate only records that match certain conditions. For example, you can migrate records that were created within a time range or split a table into multiple smaller tables.

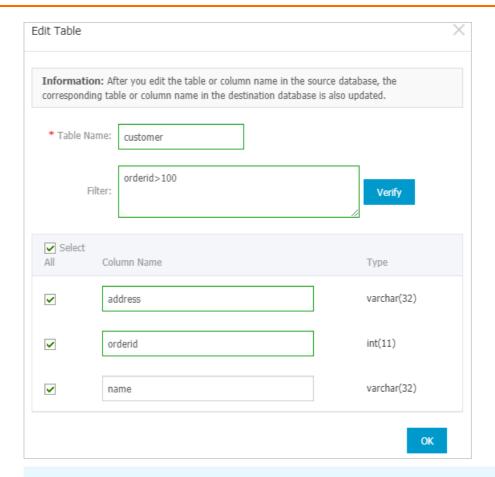
Procedure

- 1. Create a data migration task.
- 2. In the Configure Migration Types and Objects step, select the required objects into the





- 3. In the Filter text area, enter a filter expression.
 - o For relational databases, such as MySQL and SQL Server, use a SQL condition.

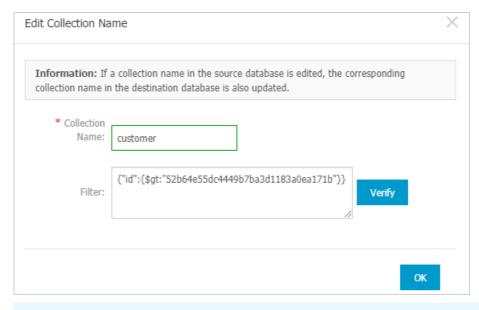


? Note

- The filter expression follows the syntax of the WHERE clause of most SQL languages.

 Do not include WHERE in the filter expression. The following operators are supported:

 = ', ' = ', ', and ' > '. For example, orderid>100'.
- You can use apostrophes (') around object names for clearer delimitation. For example, address in ('sandiego', 'tokyo') .
- For NoSQL databases, such as MongoDB, sue a JSON-formatted condition.



Note You can specify a JSON-formatted condition. For example, enter to migrate only records with an ID greater than this value.

4. Click **Verify** to verify that the specified expression is valid.

? Note

- o If the syntax is valid, you receive an Information message that the validation is passed.
- If the syntax is invalid, you receive an Error message that indicates the error in your expression. In this case, correct the filter expression and verify it again.
- 5. Click OK.
- 6. Complete other configurations and continue to create a task.

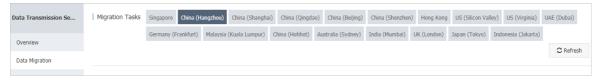
3.Throttling

3.1. Configure throttling for full data migration

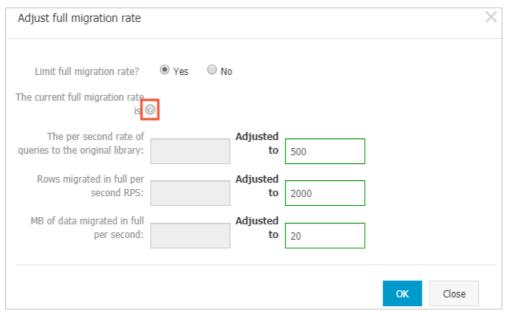
This topic describes how to modify the transmission rate of full data migration in the Data Transmission Service (DTS) console. You can set the data transmission rate to apply throttling so that the data migration workloads do not overload your source database and disrupt your normal operations.

Procedure

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



- 5. In the left-side navigation pane, click **Performance Monitoring > Full Data Migration Performance**.
- 6. In the upper-right corner of the page, click Adjust full migration rate.
- 7. Change the transmission rates.



Note You can click the question mark icon next to The current full migration rate is to view the current values of these throttling parameters.

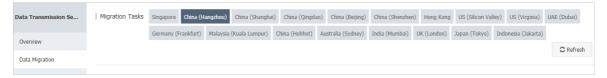
8. Click OK.

4.Migration task management 4.1. View the progress of a data migration task

This topic describes how to view the progress of a data migration task in the Data Transmission Service (DTS) console. A task defines configurations for a data replication workload, such as the connections to the source and target data stores. You can also view the task status to track the progress of your replication workloads. A data migration task consists of several phases, namely schema migration, full data migration, and incremental data migration.

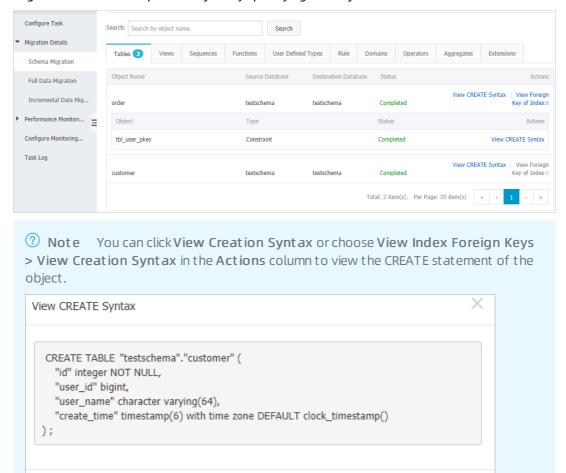
Procedure

- 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 target instance resides.

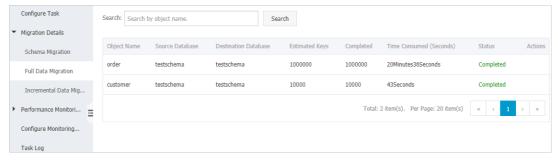


- 4. On the Migration Tasks page, click the ID of the target instance.
- 5. Depending on the migration phase, only relevant details are displayed. Perform the following operations accordingly:
 - View the details of schema migration.
 - a. In the left-side navigation pane, choose Migration Details > 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.



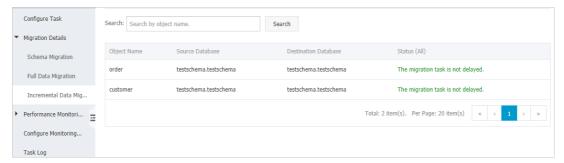
- o View the details of full data migration
 - a. In the left-side navigation pane, click Migration Details > Full Data Migration.
 - b. View the migration details of each object or search for the migration details of a specific object by specifying the object name.



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

Close

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



References

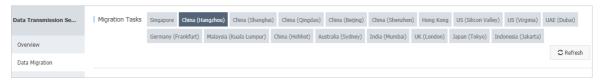
- You can check performance metrics of the full data migration phase, such as bandwidth, records per second (RPS), read/write response times, and network latency. For more information, see View the connection state and performance of full data migration
- You can check performance metrics of the full data migration phase, such as number of migrated rows, bandwidth, and data transmission rate. For more information, see View the connection status and performance of incremental data migration.

4.2. View performance metrics of full data migration

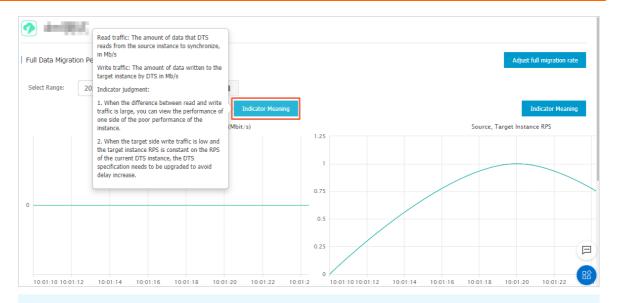
This topic describes how to view the performance metrics of the full data migration phase in the Data Transmission Service (DTS) console. DTS provides various performance metrics for the full data migration phase, including bandwidth, records per second (RPS), read/write response time, and network latency.

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. On the Migration Tasks page, click the ID of the target instance.
- 5. In the left-side navigation pane, click **Performance Monitoring > Full Data Migration Performance**.
- 6. Select a time range to display the performance profile of full data migration.



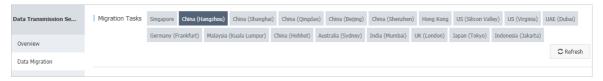
Note Move the pointer over the Indicator Meaning button in the upper-right corner of a chart to view the description of performance metrics.

4.3. View performance metrics of incremental data migration

This topic describes how to view the performance metrics of the incremental data migration phase in the Data Transmission Service (DTS) console. DTS provides various performance metrics for the incremental data migration phase, including migrated rows, bandwidth, and migration performance (QPS).

Procedure

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



- 4. On the Migration Tasks page, click the ID of the target instance.
- 5. In the left-side navigation pane, click **Performance Monitoring > Incremental Data Migration Performance**.
- 6. Select a time range to view the trend charts for the performance of incremental data migration.



Performance metrics

Performance metric	Description
Migrated rows	The total number of incremental data records that are migrated to the destination database.
Bandwidth (MBit/s)	The bandwidth between the data writing module and the data pulling module. Unit: MBit/s.
Migration performance (QPS)	The number of data records that are migrated to the destination database per second.

4.4. Restart a failed data migration task

This topic describes how to restart a failed data migration task. You can perform this operation if your data migration task is in the **Migration Failed** status during the schema migration or full data migration phase.

Restart a migration task that failed during schema migration

The task may fail during schema migration due to various causes. A typical cause is that unsupported data types are included in the migration task.

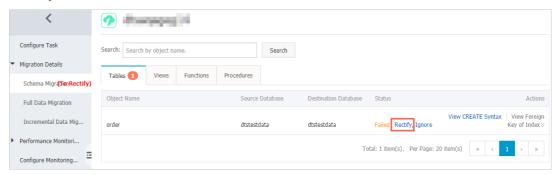
- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. At the top of the Migration Tasks page, select the region where the destination instance resides.



- 4. Use one of the following methods to resolve the issues with the failed task:
 - o Method 1
 - a. Find the target task and click View Cause and Rectify.



- b. Fix the issue based on the cause of failure that is displayed in the View Cause and Rectify message. For example, you have to modify table definitions of the source database to convert the columns into supported data types.
- c. Click Restart Task.
- o 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.



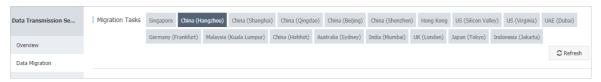
- d. Fix the issue based on the cause of failure that is displayed in the **Rectify** dialog box. For example, you have to modify table definitions of the source database to convert the columns into supported data types.
- e. Click Rectify.

? 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 all issues are resolved.
- 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 status, DTS proceeds with the data migration task, for example, entering the full data migration process.

Fix a migration task that failed 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 target 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 target database.
 - The connection to the source or target 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 target instance is insufficient or the instance is locked.
 Upgrade the target instance to a larger size or clear the log files, and then click Restart Task.
 - MylSAM tables in the source database are corrupted.
 Repair the MylSAM table in the source database, and then click Restart Task. For more information, see the official MySQL documentation.
- 5. In the dialog box that appears, fix the issue based on the cause of failure.
- 6. Click Restart Task.

4.5. Stop a data migration task

This topic describes how to stop a data migration task. If you no longer need a task or if you want to discard a failed task, you can stop the task and release the instance. After a migration task is stopped, it changes to the **Completed** status. In this status, you can only release the instance.

Prerequisites

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

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.



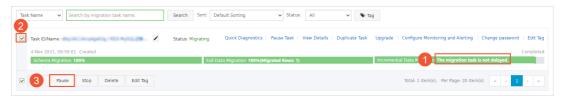
4. Perform operations based on the selected migration phases.

☐ Warning To minimize the negative impact of data migration on your normal operations, we recommend that you have a rollback plan before you switch workloads to the target instance. For more information, see Switch workloads to the target database.

Full data migration (incremental data migration not included)
 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.

target database. You must manually stop the task.

- Increment al data migration included
 A migration task with an incremental data migration phase included does not end automatically.
 Incremental data migration is an continuous process that replicates ongoing updates to the
 - a. Wait until the task progress bar shows Incremental Data Migration and The migration task is not delayed. Then, suspend write requests to the source database for a few minutes. In some cases, the progress bar shows a delay time for incremental data migration because certain newly committed updates have not been replicated.
 - b. After the status of **incremental data migration** changes to **The migration task** is **not delayed**, select the task and click **Stop** to manually stop the migration task.



5.Application switchover5.1. Switch workloads to the target database

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

Prerequisites

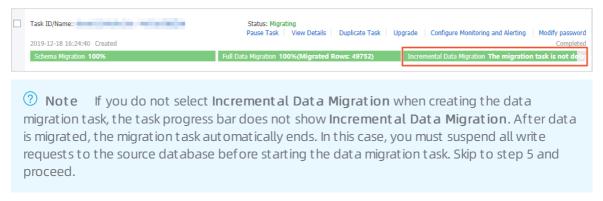
A data migration task is configured and it is in the Migrating or Completed state.

Notes

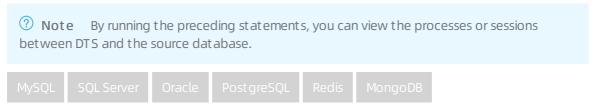
- We recommend that you switch workloads to the target database during off-peak hours to minimize the disruption to your normal operations. Before you switch workloads to the target database, you must suspend all write requests to the source database and schedule a downtime for your application.
- We recommend that you create separate accounts for data migration and assign it the required permissions. This will help you identify sessions used by DTS and manage your security more easily.

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.

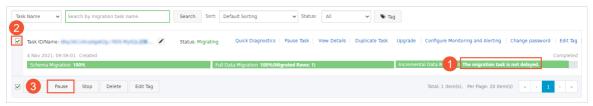


- 2. To suspend write requests to the source database, stop your application and make the source database read-only.
- 3. Log on to the source database and make sure that no new sessions are created that write into the source database. Depending on your database engine, run one of the following statements to view session information:



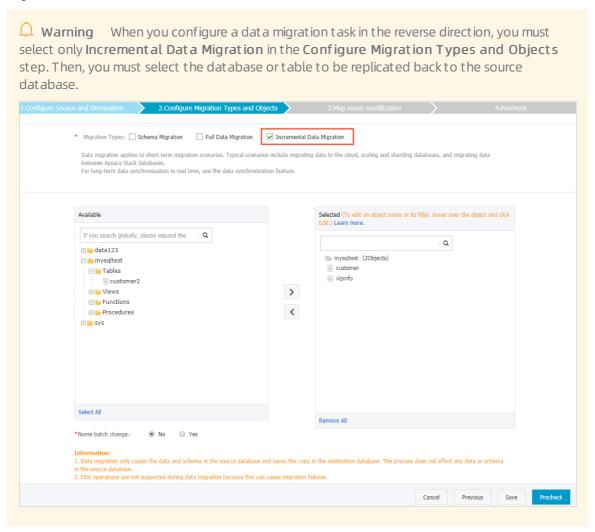
4. After the status of incremental data migration changes to The migration task is not delayed (the source and target are fully in sync), wait for one minute or longer, and then manually

stop the migration task.



- 5. While the application is stopped, make the source database writable again.
- 6. Create and start a migration task in the reverse direction so that new updates made to the target database are replicated to the source database. The migration in the reverse direction makes sure that the two databases are synchronized so that you can switch workloads back to the source database if an error occurs. When you create this migration task, select only the Incremental Data Migration phase.

For example, if migrate from a user-created MySQL database to an ApsaraDB RDS for MySQL instance, you can refer to the procedure described in Migrate data from an ApsaraDB RDS for MySQL instance to a self-managed MySQL database (select only Incremental Data Migration) to create a migration task in the reverse direction.



7. Verify that the source and target databases are in sync, modify your application to connect to the target database, and then restart your application.

35

Note With the data migration task in the reverse direction, updates made to the target database are replicated back to the source database and these two databases are in sync. If the target database fails during the test period, you can switch workloads back to the source database.

Next steps

After you switch workloads to the target database and test all application features, you can stop the task in the reverse direction (target-to-source). For more information, see Stop a data migration task.

Warning The database accounts used for data migration have the read/write permissions. After the migration is complete, you must remove permissions from these accounts or delete the accounts to avoid security risks.

FAQ

- Q: What can I do if an error occurs after I switch workloads to the target database?

 A: If an error occurs, you can switch workloads back to the source database. After you create a data migration task in the reverse direction, updates made to the target database are replicated 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 target database?
 - A: You can back up the source database before switching workloads.
- Q: What can I do if updates are unexpectedly written to the source database after I switch workloads to the target database?
 - A: You can manually compare the data objects in source and target databases and apply corrections.

6.Precheck troubleshooting6.1. Source database connectivity

Before data migration, Data Transmission Service (DTS) checks whether DTS servers can connect to the source database to avoid failures at migration time. This topic describes causes of check failure and how to fix the failure.

The connectivity check may fail due to the following reasons.

- The database account or password is incorrect
- The source instance blocks access from external IP addresses
- A firewall is enabled on the source database server
- The network is unavailable

The database account or password is incorrect

Troubleshooting:

To verify that the specified account and password are correct, you can log on to the target database from a remote database client, using the account and password.

Note If the source database is a user-created database, you can also log on locally from a client 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 source instance blocks access from external IP addresses Troubleshooting:

- If your source database is a user-created database, you can log onto the server that hosts the database system and initiate a database logon session locally from that server. If you can log on locally from the source server, it means that the source database does not allow access from external IP addresses.
- If the source database is a MySQL database, you can use a MySQL client to connect to the database and run the following statement:

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

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

Check whether the authorized IP address list includes the IP addresses of the DTS servers. For more information, see Whitelist DTS IP ranges for your user-created database.

- If the source database is an SQL Server database, check whether a firewall is enabled for the server where the database resides. You can also check whether endpoints or triggers in the source database disallow access from external IP addresses.
- If the source database is an Oracle database, check whether the TCP.VALIDNODE_CHECKING
 parameter in the sqlnet.ora configuration file is set to yes. If the parameter is set to yes, it means that the source database blocks access from external IP addresses.

Solution:

• For a source MySQL database: run the following statement to authorize the database account again:

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

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

- For a source SQL Server database: disable the triggers that block access from external IP addresses. Additionally, add DTS IP address ranges to the firewall rules or temporarily disable the firewall during migration.
- For a source Oracle database: set the TCP.VALIDNODE_CHECKING parameter to no and restart the Oracle Net Listener process.

Log on to the DTS console to perform another precheck.

A firewall is enabled on the source database server

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 enabled 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 enabled for the server.

Solution:

Add DTS IP address ranges to the firewall rules or temporarily disable the firewall during migration. Then, log on to the DTS console to perform a precheck again.

The network is unavailable

If the connectivity check failure persists after you perform the preceding troubleshooting steps, the network between DTS servers and the source database may be unavailable. You can contact Alibaba Cloud Customer Services by submitting a ticket.

6.2. Source database binary logging

If your migration task involves incremental data migration between MySQL databases, Data Transmission Service (DTS) checks the binary logging (binlog) configurations of the source database during the precheck. This topic describes how to fix issues detected by the Source Database Binary Logging check item.

Source database binary logging status

This item checks whether binary logging is enabled in the source database. Follow these instructions to fix the issue:

Solution:

- 1. Log on to the server where your user-created MySQL database is installed.
- 2. Open the my.cnf file in a text editor. In a Linux system, you can run the vim command to edit the file. Edit the text file to include the following parameters.

```
log_bin=mysql_bin
binlog_format=row
binlog_row_image=full //This parameter is required for MySQL versions later than 5.6
```

3. Restart the MySQL process. In a Linux system, run the following commands:

```
$mysql_dir/bin/mysqladmin -u root -p shutdown
$mysql_dir/bin/safe_mysqld &

Note Replace the installation directory mysql_dir of the MySQL database with the actual MySQL installation directory.
```

4. Log on to the DTS console to perform another precheck.

Source database binlog format

This item checks whether the binary log (binlog) format of the source database is ROW. Follow these instructions to fix the issue:

Solution:

- 1. Log on to the user-created MySQL database in a MySQL client.
- 2. Run the following statement to set the binlog format of the database to ROW .

```
set global binlog_format='ROW';
```

3. Restart the MySQL process. In a Linux system, run the following commands:

```
$mysql_dir/bin/mysqladmin -u root -p shutdown
$mysql_dir/bin/safe_mysqld &

Note Replace the installation directory mysql_dir of the MySQL database with the actual MySQL installation directory.
```

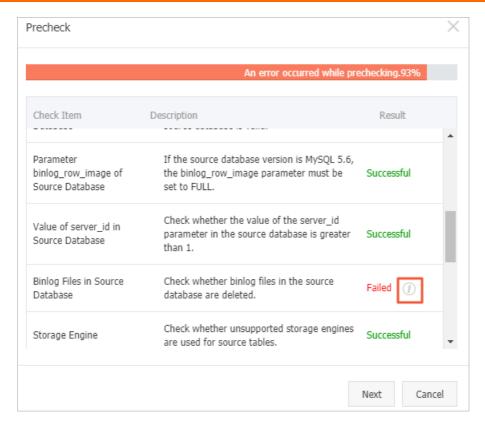
4. Log onto the DTS console to perform another precheck.

Source database binlog files

This item checks whether binary log (binlog) files of the source database are complete. If this item fails the precheck, it indicates that some binlog files are missing. Follow these instructions to fix the issue:

Solution:

1. In the **Precheck** dialog box, click the information icon next to Binlog Files in Source Database check item.



- 2. In the View More dialog box, check the deleted binlog files displayed in the Cause of Failure section. In this example, the binlog file deleted is mysql bin.000003.
- 3. Log on to the user-created MySQL database.
- 4. Run the following command on the command line to delete all binary log files before the specified log file.

```
Note In this example, the specified log file is the file right after the missing binary log file
mysql_bin.000003. Therefore, the binlog_filename is mysql_bin.000004 .

PURGE BINARY LOGS TO '<binlog_filename>';

Example:

PURGE BINARY LOGS TO 'mysql_bin.000004';
```

5. Log on to the DTS console to perform another precheck.

Parameter binlog_row_image

This item checks whether the $binlog_{row_image}$ parameter of the source database is set to Full . Before images or after images are not written to the binary log files if the item fails to pass the precheck. The following shows how to troubleshoot the problem.

Solution:

- 1. Log on to the user-created MySQL database.
- 2. Run the following command on the command line and set binlog row image to FULL .

```
set global binlog_row_image=FULL;
```

3. Log on to the DTS console to perform another precheck.

6.3. Integrity constraints

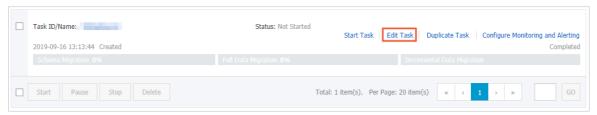
Before data migration, Data Transmission Service (DTS) checks the integrity constraints in the source database to make sure that the integrity constraints still hold after the migration.

This item checks whether all the parent and child tables with foreign key dependencies with the selected objects are also included in the migration scope. If this precheck item fails, certain parent tables are not selected in the migration scope. You can choose one of the following methods to fix this issue:

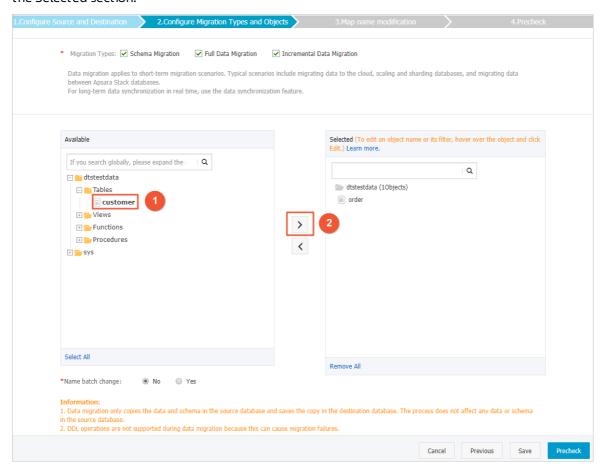
Add parent tables to the migration scope

Follow these steps to include parent tables of the tables that fail the precheck.

1. Find the migration task and click Edit Task.



- 2. Click Set Whitelist and Next in the lower-right corner of the page.
- 3. Select parent tables of the tables that fail to pass the precheck and click to move them to the Selected section.

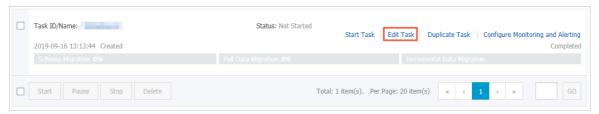


4. Click Precheck.

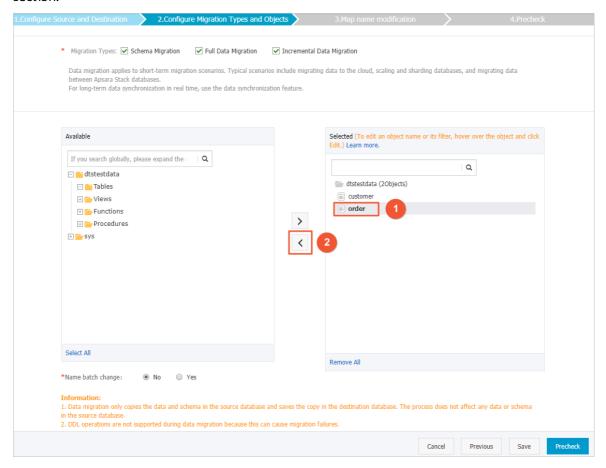
Remove child tables from the migration scope

Follow these steps to remove tables that fail the precheck from the objects to be migrated.

1. Find the migration task and click **Edit Task**.



- 2. Click Set Whitelist and Next in the lower-right corner of the page.
- 3. Select tables that fail to pass the precheck and click to remove them from the Selected section.



4. Click Precheck.

Delete foreign key dependencies

Follow these steps to delete foreign key dependencies from the source tables that fail to pass the precheck.

- 1. Log on to the source database.
- 2. Delete foreign key dependencies from the tables that fail the precheck.

Note If the source database is on a managed database service provided by Alibaba Cloud, you can also log on to the DMS console to perform this operation. For more information, see Manage foreign keys.

3. Perform another precheck.

6.4. MySQL storage engine

If your migration task involves incremental data migration between MySQL databases, Data Transmission Service (DTS) checks storage engines of the source database during the precheck. This topic describes how to fix issues detected by the Storage Engine check item.

This item checks whether storage engines used in the source database support incremental data migration. For example, the FEDERATED and MRG_MYISAM storage engines do not support incremental data migration. If this item fails to pass the precheck, the source database may have one of the unsupported storage engines.

Solution

- 1. In the precheck results report, click the information icon next to the error item.
- 2. In the View More dialog box that appears, check the table name displayed in the Cause of Failure section.
- 3. Find the data migration task and click Edit Task.



- 4. Click Set Whitelist and Next in the lower-right corner of the page.
- 5. Select the tables whose storage engines fail to pass the precheck and click to move them out of the Selected section.
- 6. Click Precheck.

More information

To migrate tables with the FEDERATED and MRG_MYISAM storage engines, you can create a data migration task for these tables and select **Full Data Migration** and **Schema Migration** only.

6.5. Source database permissions

Before data migration, Data Transmission Service (DTS) checks permissions granted to the source database account to make sure that the specified account is authorized to perform the required migration operations.

The required permissions vary, depending on the source database and the migration phases involved. Refer to the documentation about the prerequisites for each migration procedure. For more information, see Data migration procedure list

6.6. Target database permissions

Before data migration, Data Transmission Service (DTS) checks permissions granted to the target database account to make sure the specified account is authorized to perform the required migration operations.

The required permissions vary, depending on the target database and the migration phases involved. Refer to the documentation about the prerequisites for each migration procedure. For more information, see Data migration procedure list

6.7. Object name conflicts

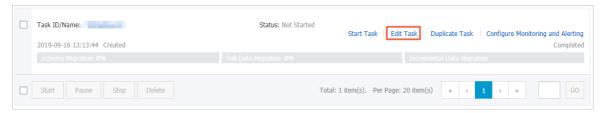
Before data migration, Data Transmission Service (DTS) checks whether the objects to be migrated already exist in the target database. This topic describes how to find the same-named objects in the target database and rename these objects.

This check item fails if objects in the target database have identical names with the objects to be migrated. To resolve this issue, use one of the following methods:

Configure object name mappings

You can configure name mappings to rename the objects with conflicting names for them to have different names in the target database. To do this, follow these steps:

1. Find the migration task and click Edit Task.



- 2. Click Set Whitelist and Next in the lower-right corner of the page.
- 3. Rename objects migrated to the target database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping.
- 4. Click Precheck.

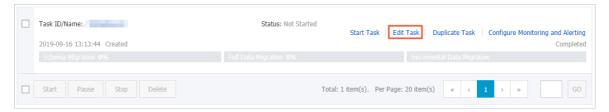
Rename objects in the target database

Log on to the target database, and rename or delete the objects that have identical names with the objects to be migrated.

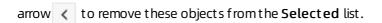
Remove objects with conflicting names from the migration scope

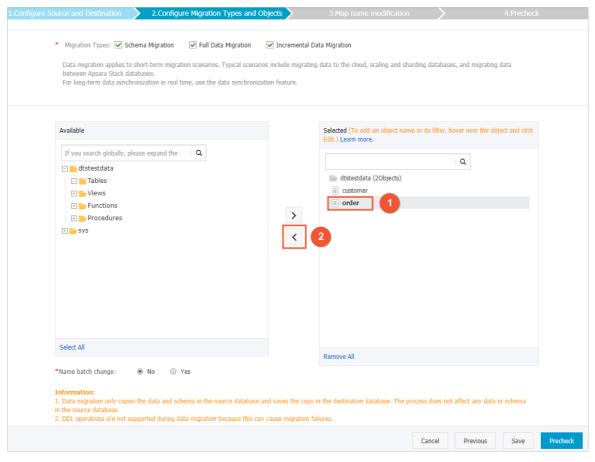
Modify the task configurations to exclude the objects that have conflicting names from the migration scope.

1. Find the migration task and click Edit Task.



- 2. Click Set Whitelist and Next in the lower-right corner of the page.
- 3. Select the objects that have identical names with those in the target database and click the left





4. Click Precheck.

6.8. Target database availability

Before data migration, Data Transmission Service (DTS) checks whether a database with the same name as the source database exists in the target database system. If no such a database exists in the target instance, DTS creates a database. If the automatic creation fails, the precheck fails. In the following cases, DTS cannot automatically create a database in the target instance and the precheck fails.

The name or character set of the source database do not comply with rules in the target RDS instance.

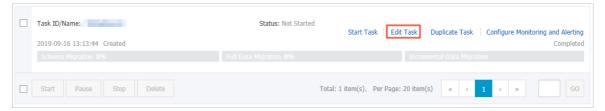
- The name of the source database contains special characters other than underscores (_) and hyphens (-).
- The character set of the source database is not UTF-8, GBK, Latin1, or UTF-8 (MB4).

Solution:

1. Create a database in the target RDS instance that complies with the naming convention and character set rules. For more information, see Create a database for an RDS for MySQL instance.

Note When you create the database, you must grant permissions to the target database account that is used by the data migration task.

- 2. Log on to the DTS console.
- 3. In the left-side navigation pane, click **Dat a Migration**.
- 4. Find the data migration task and click Edit Task.



- 5. Click Set Whitelist and Next in the lower-right corner of the page.
- 6. Configure the name of the target database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping.
- 7. Click Precheck.

The permissions granted to the target database account are insufficient

Solution:

- Grant read and write permissions to the target database account. For more information about how
 to change permissions, see Modify the permissions of a standard account on an ApsaraDB RDS for
 MySQL instance
- 2. Perform another precheck

6.9. server_id in the source database

Before you start incremental data migration between MySQL databases, DTS checks the value of the server-id parameter in the source database. This topic describes how to configure the server-id parameter to pass the precheck.

Solution

- 1. Log on to the source user-created MySQL database.
- 2. Run the following statement to view the current server id value:

```
show variables like '%server_id%';
```

3. Run the following statement to set the server-id parameter.

```
set global server id=<id>;
```

- Note Replace the <id> placeholder with a valid numeric server ID. Follow these rules when determining a server ID:
 - The server ID must be an integer greater than 1.
 - If your source database is in a primary/secondary replication deployment, make sure that the assigned server ID is unique in the cluster of your source database system.
 - The server ID is overwritten after you restart the database service. To make the setting permanent, apply the change into the corresponding configuration file.

Example:

```
set global server_id=2;
```

4. Log on to the DTS console and perform another precheck.

6.10. Target database connectivity

Before data migration, Data Transmission Service (DTS) tests the connectivity between DTS servers and the target database. This topic lists common causes of check failures and how to troubleshoot these failures.

The connectivity check may fail due to the following reasons.

- The database account or password is incorrect
- The target database blocks access from external IP addresses
- A firewall is enabled on the target database server
- The network is unavailable

The database account or password is incorrect

Troubleshooting:

To verify that the specified account and password are correct, you can log on to the target database from a remote database client, using the account and password.

Solution:

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

The target database blocks access from external IP addresses Troubleshooting:

- If your target database is a user-created database, you can log on to the server that hosts the database system and initiate a database logon session locally from that server. If you can log on locally from the target server, it means that the target database does not allow access from external IP addresses.
- If the target database is a MySQL database, you can use a MySQL client to connect to the database and run the following statement:

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

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

Check whether the authorized IP address list includes the IP address ranges of DTS servers. For more information, see Whitelist DTS IP ranges for your user-created database.

- If the target database is an SQL Server database, check whether a firewall is enabled for the server where the database resides. You can also check whether endpoints or triggers in the target database disallow access from external IP addresses.
- If the target database is an Oracle database, check whether the TCP.VALIDNODE_CHECKING parameter in the sqlnet.ora configuration file is set to yes. If the paremeter is set to yes, the target database blocks access from external IP addresses.

Solution:

• For a target MySQL database: run the following statement to authorize the database account again:

```
@ Note Replace the username and password placeholders with the database account and password specified for the data migration task.
```

- For a target SQL Server database: disable the triggers that block access from external IP addresses. Additionally, add DTS IP address ranges to the firewall rules or temporarily disable the firewall during migration.
- For a target Oracle database: set the TCP.VALIDNODE_CHECKING parameter to no and restart the Oracle Net Listener process.

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

A firewall is enabled on the target database server

Troubleshooting:

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

Solution:

Add DTS IP address ranges to the firewall rules or temporarily disable the firewall during migration. Then, log on to the DTS console to perform a precheck again.

The network is unavailable

If the connectivity check failure persists after you perform the preceding troubleshooting steps, the network between DTS servers and the source database may be unavailable. You can contact Alibaba Cloud Customer Services by submitting a ticket.

6.11. Source database engine version

Before data migration, Data Transmission Service (DTS) checks the version number of the source database engine to make sure that the selected source database engine is supported. For more information about supported sources and targets, see Supported sources and targets.

7.Migrate from a user-created database to Alibaba Cloud

7.1. Source: MySQL

7.1.1. Migrate data from a user-created MySQL database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a user-created 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 user-created MySQL database, you can select 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 user-created 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 user-created MySQL database.

Precautions

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the database load. 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 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 usage 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, duplicate data may exist in the destination database.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the float or double data type. If the precision is not specified, 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 how to create a database and the database naming conventions, see Create databases and accounts.

• DTS automatically resumes a failed data migration 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 from Alibaba Cloud over the Internet. For more
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.



- During schema migration, DTS changes the value of the SECURITY attribute in views, stored procedures, and functions from DEFINER to INVOKER.
- DTS does not migrate user information. Before a user can call views, stored procedures, and functions of the destination database, you must grant the read/write permissions to the user.
- Full data migration

DTS migrates historical data of the required objects from the user-created 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 instance. After full data migration is complete, the tablespace of the destination instance is larger than that of the source database.

• Increment al dat a migration

After full data migration is complete, DTS retrieves binary log files from the user-created MySQL database. Then, DTS synchronizes incremental data from the user-created 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 user-created MySQL database to Alibaba Cloud.

SQL operations that can be synchronized during incremental data migration

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE

Operation type	SQL statements
DDL	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
User-created 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/write permissions	The read/write permissions	The read/write permissions

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

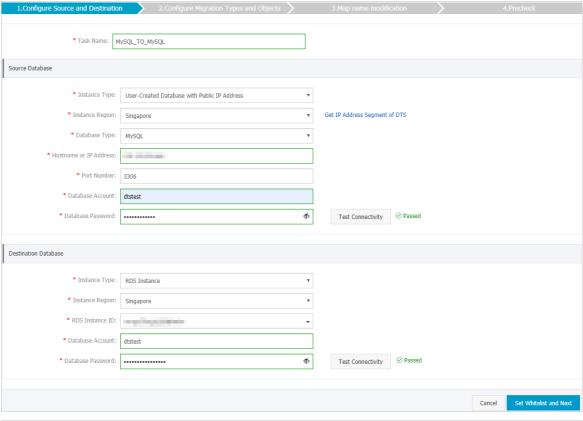
- Create an account for a user-created MySQL database and configure binary logging for a user-created MySQL database
- Create an account for an RDS for MySQL instance and Change the permissions of an account for an RDS for MySQL instance

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. 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 for the data migration task.



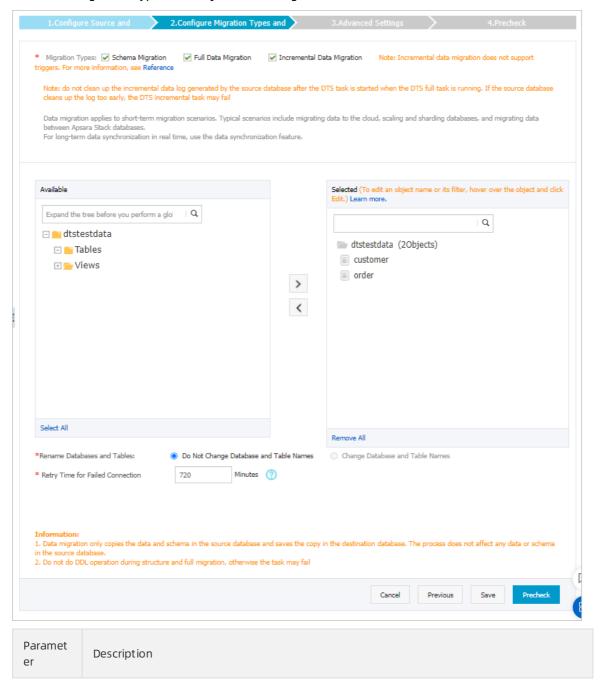
Section	Parameter	Description	
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.	
		Select an instance type based on where the source database is deployed. The procedure in this topic uses User-Created Database with Public IP Address as an example.	
	Instance Type	Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview.	
		If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.	
Instance Region		Note If a whitelist is configured for the user-created MySQL database, you must manually add the CIDR blocks of DTS servers to the whitelist of the user-created MySQL 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			

Section	Parameter	Description
	Hostname or IP Address	Enter the IP address that is used to access the user-created MySQL database. In this example, enter the public IP address.
	Port Number	Enter the service port number of the user-created MySQL database. The default port number is 3306 .
	Dat abase Account	Enter the account of the user-created MySQL database. For more information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password for the source database account.
	Database Password	Note After you specify the information about the selfmanaged 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 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 more information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password for the destination database account.
Destinatio n Database	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.
	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 for an RDS for MySQL instance.
		Note The Encryption parameter is available only for regions in mainland China and the Hong Kong (China) region.

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

Note The CIDR blocks of DTS servers are automatically added 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 objects to be migrated.



Objects O To perform only full data migration, select Schema Migration and Full Data Migration. To migrate data with minimal downtime, select Schema Migration, Full Data Migration, and Incremental Data Migration. O Note If Incremental Data Migration is not selected, do not write data into the source database during full data migration. This ensures data consistency between the source and destination databases. Select objects from the Available section and click the picon to move the objects to the Selected section. Note You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains the same as that in the source database. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.	Paramet er	Description
Note If Incremental Data Migration is not selected, do not write data into the source database during full data migration. This ensures data consistency between the source and destination databases. Select objects from the Available section and click the icon to move the objects to the Selected section. Note You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains the same as that in the source database. You can change the names of the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that	Migratio	Migration. o To migrate data with minimal downtime, select Schema Migration, Full Data
the Selected section. Note You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains the same as that in the source database. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that	n Types	the source database during full data migration. This ensures data consistency
 You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains the same as that in the source database. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that 		
	Objects	 You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains the same as that in the source database. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that

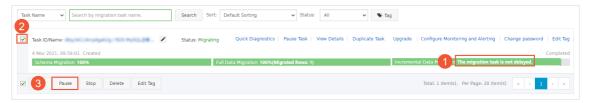
- 8.
- 9. After the data migration task passes the precheck, click **Next**.
- 10. In the Confirm Settings dialog box, configure the Channel Specification parameter. Then, read and select Data Transmission Service (Pay-as-you-go) Service Terms.
- 11. Click **Buy and Start** to start the migration task.

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 stop the migration task by using the following procedure.

- Full dat a 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 to do next

The database accounts used for data migration have the read/write permissions. After data is migrated, you must delete the database accounts to ensure security.

7.1.2. Migrate data from a user-created MySQL database connected over Express Connect, VPN Gateway, or Smart Access Gateway to an ApsaraDB RDS for MySQL database

This topic describes how to migrate data from a user-created 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. To migrate data from a user-created MySQL database, you can select all of the supported migration types to ensure service continuity.

Prerequisites

- The version of the user-created 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 database is larger than the total space of the data in the user-created MySQL database.
- The on-premises network to which the user-created MySQL database belongs is connected to Alibaba Cloud over Express Connect, VPN Gateway, or Smart Access Gateway.



Note For more information, see Connect an on-premises database to Alibaba Cloud.

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
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.



- 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.

• Increment al dat a 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

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:

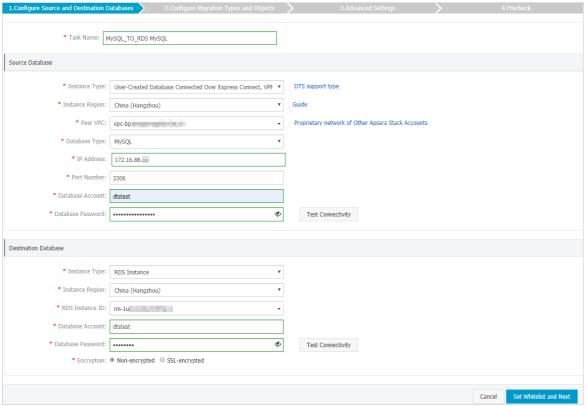
- Self-managed MySQL database: Create an account for a user-created MySQL database and configure binary logging
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

Preparations

- 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.
- 2.
- 3.
- 4.
- 5. Configure the information about the source and destination databases for the data migration task.



	5
Parameter	Description
Task Name	DTS automatically generates a task name. We recommend that you use 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 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 connect to the user-created MySQL database.
Port Number	Enter the service port number of the user-created MySQL database. The default port number is 3306.
Dat abase Account	Enter the account for the user-created MySQL database. For more information about permissions required for the account, see Permissions required for database accounts.
	Instance Type Instance Region Peer VPC Database Type IP Address Port Number Database

Section	Parameter	Description	
		Enter the password for the database account.	
	Dat abase Password	? Note	
	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.	
	Database Account	Enter the database account of the destination RDS instance. For more information about permissions required for the account, see Permissions required for database accounts.	
Destinatio		Enter the password for the database account.	
Database Password	Dat abase Password	? Note	
	Encryption	Select Non-encrypted or SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the RDS instance before configuring the data migration task. For more information, see Configure SSL encryption for an RDS for MySQL instance.	
		Note The Encryption parameter is available only in mainland China and Hong Kong(China).	

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

? Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination ApsaraDB RDS for MySQL instance. This ensures that DTS servers can connect to the destination RDS instance.

7.

8.

9.

10.

11.

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

7.1.3. Migrate data from a user-created 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 user-created 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.

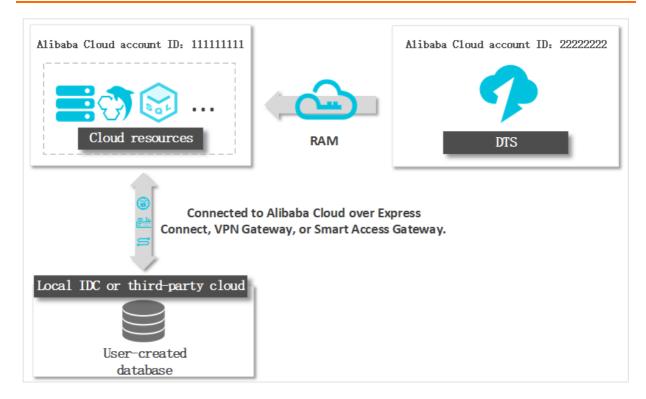
Prerequisites

- The version of the user-created 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 user-created MySQL database.
- The on-premises network to which the user-created 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.

Note For more information about how to connect an on-premises data center to a VPC, see Connect an on-premises database to Alibaba Cloud.

Context

The on-premises data center is connected to Alibaba Cloud VPC over Express Connect, VPN Gateway, or Smart Access Gateway. You need to migrate data from a user-created MySQL database that resides in an on-premises data center to an RDS 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 Alibaba Cloud over the Internet. For more
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.



- o During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- o 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 self-managed MySQL database to the destination database in the ApsaraDB RDS for MySQL instance.

② 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 dat abase. 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	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

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:

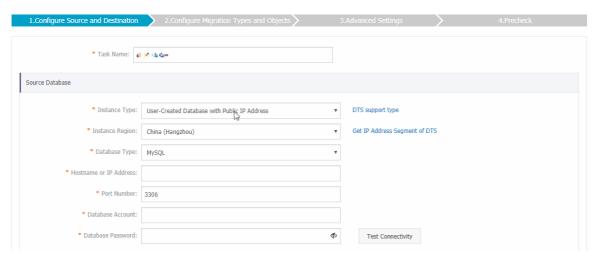
- Self-managed MySQL database: Create an account for a user-created MySQL database and configure binary logging
- ApsaraDB RDS for MySQL instance: Create an account on an ApsaraDB RDS for MySQL instance and Modify the permissions of a standard account on an ApsaraDB RDS for MySQL instance.

Preparations

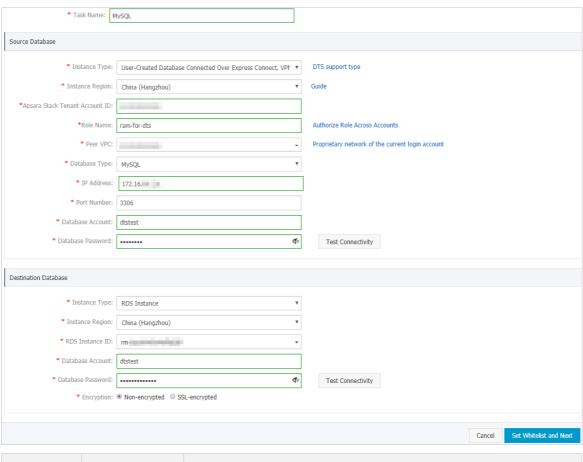
- 1. Create an account for a user-created MySQL database and configure binary logging.
- 2. Log on to the Alibaba Cloud console by using the Alibaba Cloud account to which the Express Connect circuit belongs. Authorize DTS to access the network that is connected over Express Connect. For more information, see Configure a route between DTS and Express Connect, VPN Gateway, or Smart Access Gateway.
- 3. Create a RAM role and authorize this role to access the cloud resources of the Alibaba Cloud account to which the Express Connect circuit belongs. 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. Log on to the DTS console by using the Alibaba Cloud account to which the destination RDS instance belongs.
- 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.



6. Configure the source and destination databases for the data migration task.



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.

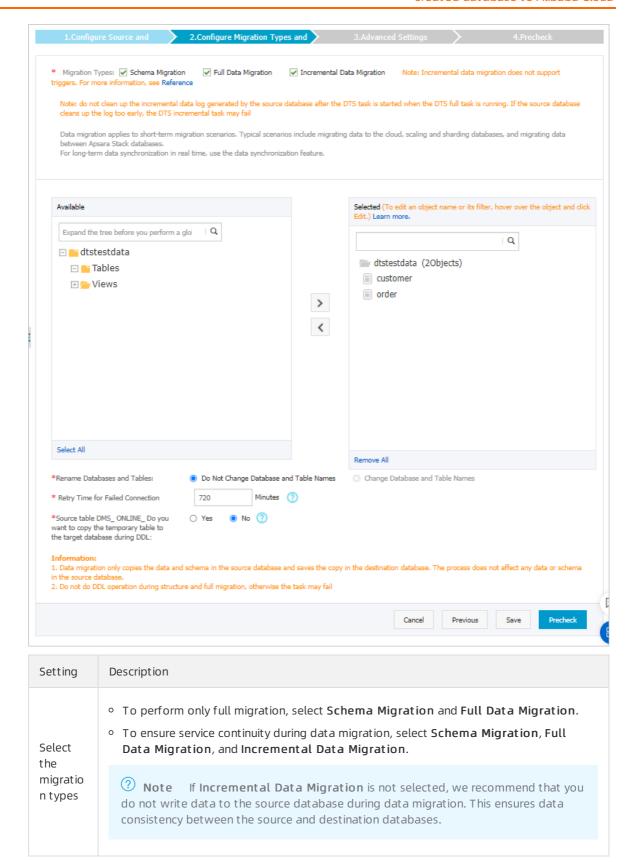
Section	Parameter	Description		
	Instance Region	Select the region of the VPC that is connected to the user-created MySQL database.		
	Alibaba Cloud Account ID	Enter the ID of the Alibaba Cloud account to which the Express Connect circuit belongs.		
		Note To obtain the ID of the Alibaba Cloud account to which the Express Connect circuit belongs, 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 Logon Account: do ""@first allyund com Edit Verified: No real-name authentication Go to certification Security Settings Real-name Verification Three-party binding: No bending: Data! Contact Management C		
Source Database	Role Name			
Peer	Peer VPC	Select the ID of the VPC that is connected to the user-created MySQL database.		
Database Type Select MySQL .		Select MySQL.		
	IP Address	Enter the endpoint that is used to access the user-created MySQL database.		
Database Account Database Account		Enter the service port number of the user-created MySQL database. The default port number is 3306 .		
		Enter the account of the user-created 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 source 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 RDS Instance.		
	Instance Region	Select the region where the destination RDS instance resides.		
RDS Instance ID Select the ID of the destination RDS instance.		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.	
Destinatio		Enter the password of the destination database account.	
n Database	Database 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.	
	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 for an RDS MySQL instance.	
	Епстурстоп	Note The Encryption parameter is available only for regions in mainland China and the Hong Kong (China) region.	

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 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.



Setting	Description
	Select one or more objects from the Available section and click the icon to move the objects to the Selected section.
Select the objects to be migrate d	 Note You can select columns, tables, or databases as the objects to be migrated. 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. 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 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	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.
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	Note If online DDL operations generate a large amount of data, the data migration task may be delayed.
	 No: 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.
	Note If you select No, the tables in the destination database may be locked.

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



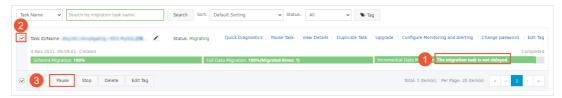
- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- \circ If the task fails to pass the precheck, you can click the \bigcirc icon next to each failed item

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.
- 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.



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

7.1.4. Migrate data from a user-created MySQL database to an Apsara PolarDB for MySQL cluster

Apsara PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. The service 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 user-created MySQL database to an Apsara PolarDB for MySQL cluster by using Data Transmission Service (DTS).

Prerequisites

- The version of the user-created MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- The destination Apsara PolarDB for MySQL cluster is created. For more information, see Create a PolarDB MySQL cluster.
- If the source MySQL database is an on-premises database, the CIDR blocks of DTS servers are added to the IP whitelist of the source MySQL 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. 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.

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 Alibaba Cloud over the Internet. For more
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	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

Permissions required for database accounts

Database	Schema migration and full data migration	Incremental data migration
User-created MySQL database	The SELECT permission	The REPLICATION CLIENT, REPLICATION SLAVE, SHOW VIEW, and SELECT permissions
Apsara PolarDB for MySQL cluster	The read/write permissions	The read/write permissions

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

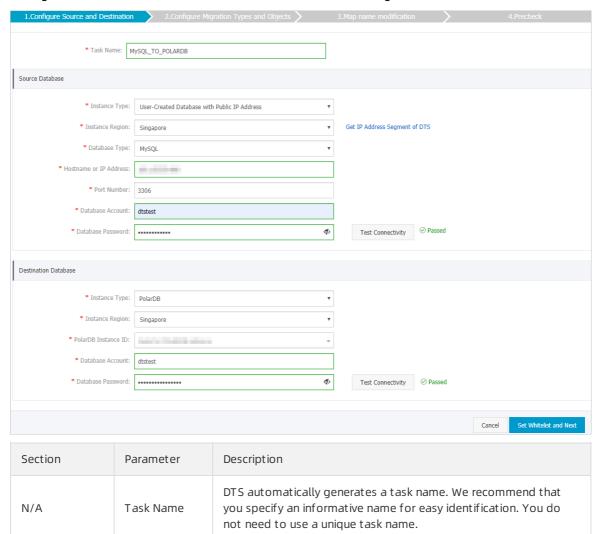
- Create an account for a user-created MySQL database and configure binary logging for a user-created MySQL database
- Create a database account for an Apsara PolarDB for MySQL cluster

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 for the data migration task.



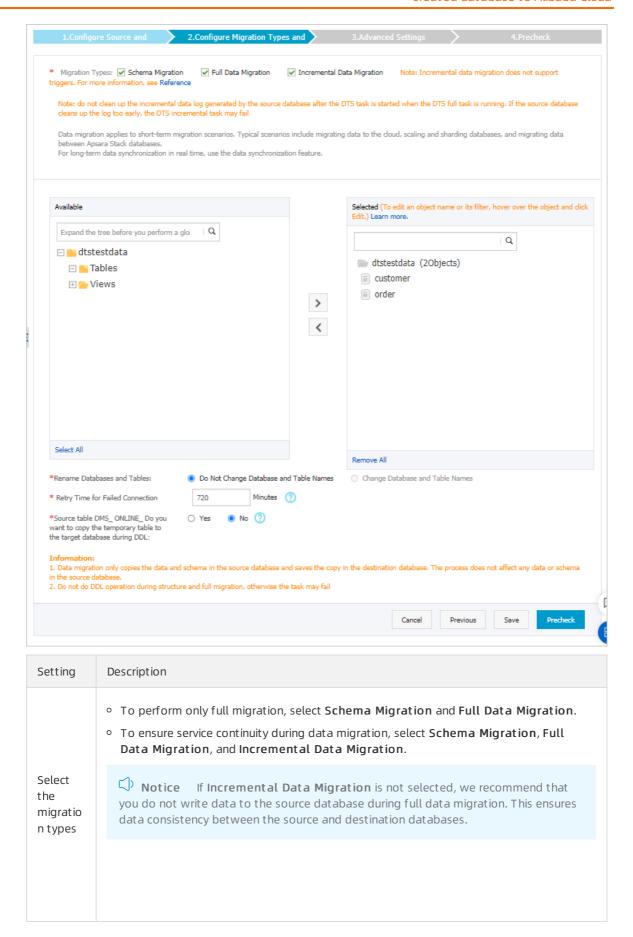
Section	Parameter	Description	
	Instance Type	Select an instance type based on where the source database is deployed. In this example, select User-Created Database with Public IP Address . Note If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see Preparation overview.	
	Instance Region	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region. ② Note If a whitelist is configured for the user-created 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.	
	Dat abase Type	Select MySQL.	
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created MySQL database. In this example, enter the public IP address.	
	Port Number	Enter the service port number of the user-created MySQL database. The default port number is 3306 .	
	Dat abase Account	Enter the account of the user-created 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 source database account.	
Source Database		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 PolarDB.	
	Instance Region	Select the region where the destination PolarDB cluster resides.	
	PolarDB Instance ID	Select the ID of the destination PolarDB cluster.	
Destination	Database Account	Enter 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		Enter the password of the destination 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 destination Apsara 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.



Setting	Description
	Select one or more objects from the Available section and click the icon to move the objects to the Selected section.
Select the objects to be migrate d	 Notice You can select columns, tables, or databases as the objects to be migrated. 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. 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 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	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.
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	Note If online DDL operations generate a large amount of data, the data migration task may be delayed.
perform s online	 No: 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	Note If you select No, the tables in the destination database may be locked.
the source table	

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

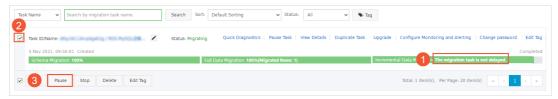


- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- \circ If the task fails to pass the precheck, you can click the \bigcirc icon next to each failed item

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.
- 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.

7.1.5. Migrate data from a user-created MySQL database to a DRDS instance

This topic describes how to migrate data from a user-created MySQL database to a Distributed Relational Database System (DRDS) instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When you migrate data from a user-created MySQL database to a DRDS instance, you can select all of the supported migration types to ensure service continuity.

Prerequisites

- The version of the user-created MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- A database is created in the DRDS instance based on one or more ApsaraDB RDS for MySQL instances
 that you purchased. DTS does not support databases that are created based on Apsara PolarDB for
 MySQL clusters.
- The available storage space of the destination RDS instance is larger than the total size of the data in the user-created MySQL database.

Precautions

79

- DTS does not support schema migration from a user-created MySQL database to a DRDS instance.
 - **Note** During schema migration, DTS migrates the schemas of the required objects such as tables 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 database will overwrite the data in the destination instance after the task is resumed.

Billing

Migration type	Instance configuration	Internet traffic
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 user-created MySQL database to the destination DRDS 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 dat a migration

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

SQL operations that can be synchronized during incremental data migration

INSERT, UPDATE, DELETE, and REPLACE

Permissions required for database accounts

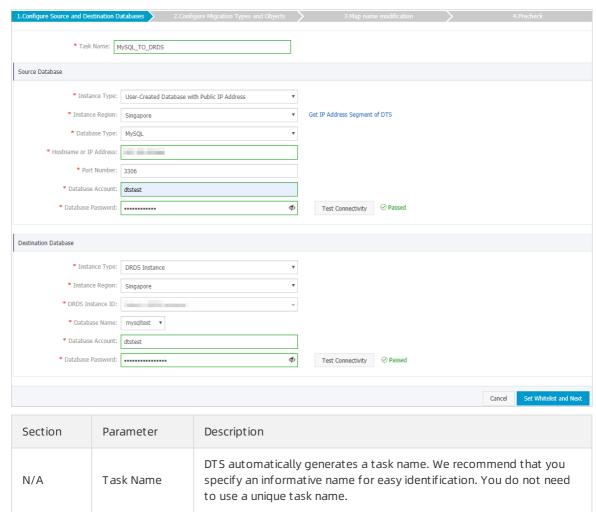
Database	Full data migration	Incremental data migration
User-created MySQL database	The SELECT permission	The REPLICATION CLIENT, REPLICATION SLAVE, SHOW VIEW, and SELECT permissions
DRDS instance	The read/write permissions	The read/write permissions

Before you begin

- Create an account for the user-created MySQL database and enable binary logging. For more information, see Create an account for a user-created MySQL database and configure binary logging.
- 2. Create a database and tables in the destination DRDS instance based on the schema to be migrated from the user-created MySQL database. For more information, see Create a DRDS database and Create a 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 for the data migration task.



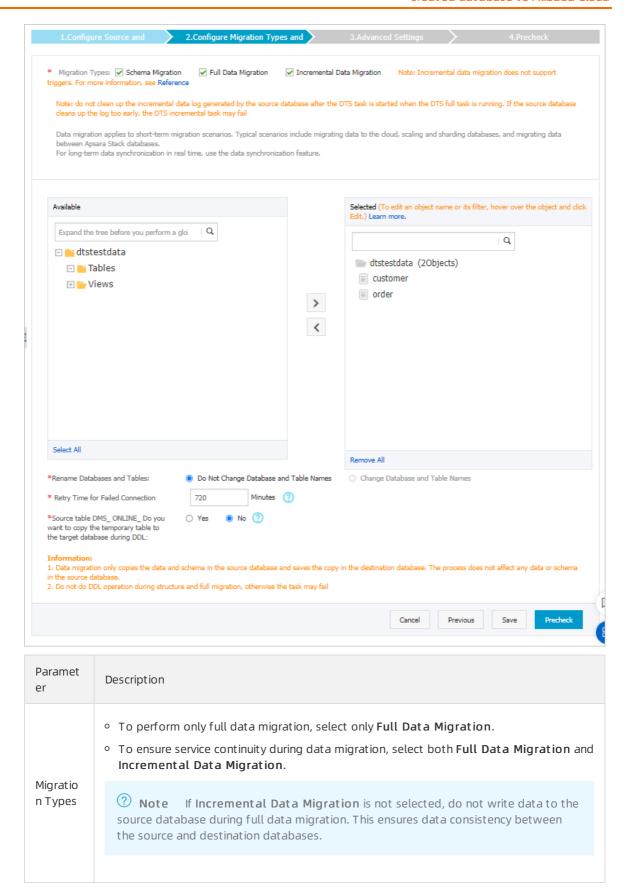
Section	Parameter	Description
	Instance Type	Select an instance type based on where the source database is deployed. In this example, select User-Created Database with Public IP Address .
		Note If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see Preparation overview.
		If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.
	Instance Region	Note If a whitelist is configured for the user-created 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.
Source		
Database	Database Type	Select MySQL.
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created MySQL database. In this example, enter the public IP address.
	Port Number	Enter the service port number of the user-created MySQL database. The default port number is 3306.
	Dat abase Account	Enter the account of the user-created 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 source database account.
	Dat abase Password	Note After you specify the information about the selfmanaged 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	Select the region where the destination DRDS instance resides.
	DRDS Instance	Select the ID of the destination DRDS instance.

Section	Parameter	Description	
Destinatio n	Dat abase Name	Select the name of the destination database.	
Database	Dat abase Account	Enter the database account of the destination DRDS 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 destination database account.	
		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.

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

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

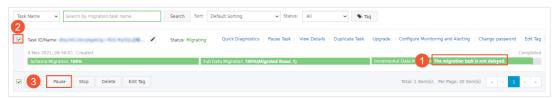


Paramet er	Description
	Select objects from the Available section and click the icon to move the objects to the Selected section.
Objects	 Note You can select columns, tables, or databases as the objects to be migrated. If you select tables or columns as the objects to be migrated, DTS does not migrate other objects such as views, triggers, and stored procedures to the destination database. 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. 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.

8.

- 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 self-managed 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 DRDS instance.

What to do next

The database accounts used for data migration have the read/write permissions. After data migration is complete, you must delete the accounts of both the user-created MySQL database and the database in the DRDS instance.

7.2. Source: SQL Server

7.2.1. Migrate incremental data from a user-created SQL Server database to an ApsaraDB RDS for SQL Server instance

Data Transmission Service (DTS) provides an incremental data migration phase that allows you to perform a minimized-downtime migration. You do not have to suspend write requests to the source database during migration. The migration task captures ongoing changes that occur during the full data migration and reapplies the changes after the full data migration is complete.

Prerequisites

• The version of the user-created SQL Server database is 2008, 2008 R2, 2012, 2014, 2016, or 2017.



- o SQL Server Cluster or SQL Server AlwaysOn High Availability Group is not supported.
- If you migrate data between different versions of databases, make sure that the database versions are compatible.
- The tables to be migrated from the user-created SQL Server database must 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 user-created 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 user-created 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 how to create a database and the database naming conventions, see Create a database for an RDS SOL 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, and 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 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, and default.
- Full data migration
 DTS migrates historical data of the required objects from the user-created SQL Server database to the destination ApsaraDB RDS for SQL Server instance.
- Increment al data migration
 After full data migration is complete, DTS synchronizes increment al data from the user-created SQL
 Server database to the destination ApsaraDB RDS for SQL Server instance. Increment al data migration allows you to ensure service continuity when you migrate data from a user-created SQL Server database to Alibaba Cloud.

SQL operations that can be synchronized during incremental data migration

- INSERT, UPDATE, and DELETE
 - Note The UPDATE operations that update only the large fields cannot be synchronized.
- CREATE TABLE
 - **? Note** The CREATE TABLE operations used to create partition tables or tables that contain functions cannot be synchronized.
- ALTER TABLE operations, only including 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
User-created SQL Server database	The SELECT permission	The SELECT permission	The sysadmin permission
ApsaraDB RDS for SQL Server instance	The read/write permissions	The read/write permissions	The read/write permissions

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

- CREATE USER for a user-created SQL Server database
- Create an account for an RDS for SQL Server instance for an ApsaraDB RDS for SQL Server instance

Incremental data migration

To avoid data migration failures caused by dependencies between 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.

4. Perform increment al data migration.

Note 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.

Before you begin

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

1. Run the following command in the user-created SQL Server database to change the recovery mode 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 **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 5 8 1

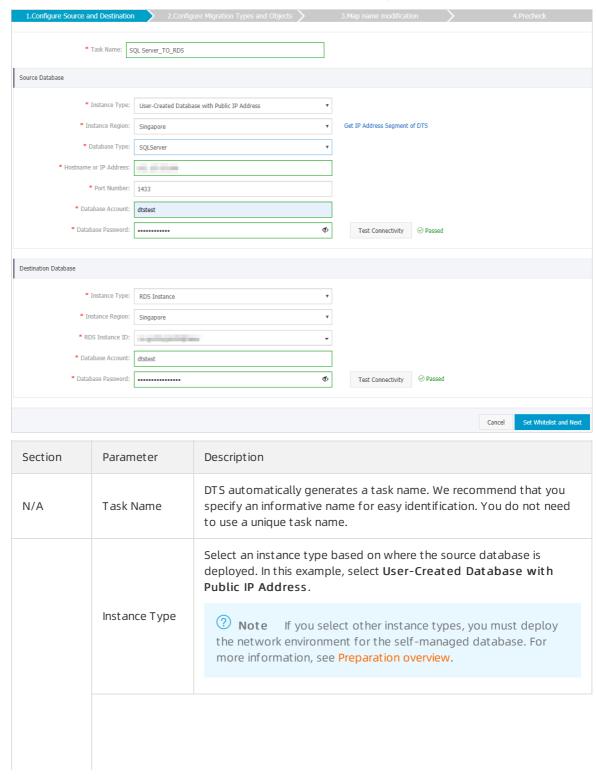
- <database_name>: the name of the source database.
- <physical_backup_device_name>: the storage path and file name of the backup file.

Example

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

Procedure

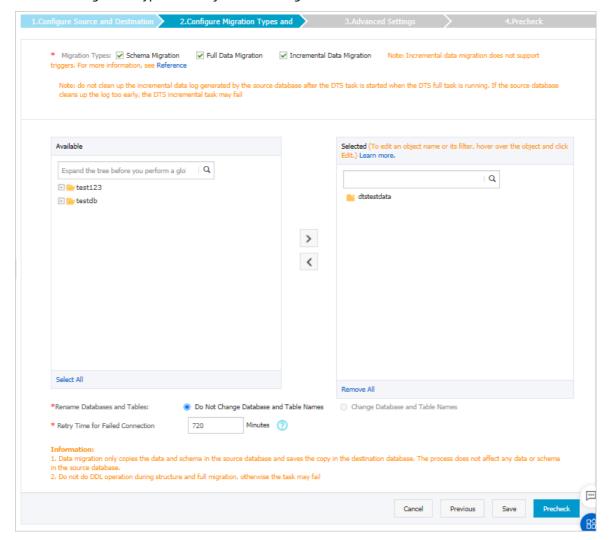
- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the upper part of the Migration Tasks page, select the region where the RDS instance resides.
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and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



Section	Parameter	Description		
Source	Instance Region	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region. Note If a whitelist is configured for the user-created 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	Dat abase Type	Select SQL Server.		
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created SQL Server database. In this example, enter the public IP address.		
	Port Number	Enter the service port number of the user-created SQL Server database. The default port number is 1433. Note The service port of the user-created SQL Server database must be accessible over the Internet.		
	Dat abase Account	Enter the account of the user-created SQL Server database.		
	Dat abase Password	Phote After you specify the information about the selfmanaged 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 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.		
Destinatio n Database				

Section Parameter	Description		
	Enter the password of the destination 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.
 - **? 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 objects to be migrated.

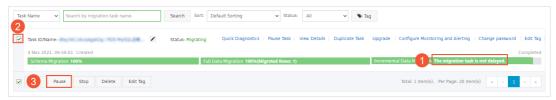


Paramet er	Description
Migratio n Types	 To perform only full data migration, select Schema Migration and Full Data Migration. To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration. In this example, select all of the three migration types.
ITTypes	Note If Incremental Data Migration is not selected, do not write data to the user-created SQL Server database during full data migration. This ensures data consistency between the source and destination databases.
	Select objects from the Available section and click the joicon to move the objects to the Selected section.
Objects	 Note 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. You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains unchanged. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.

8.

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

7.2.2. Migrate full data from a user-created SQL Server database to ApsaraDB RDS for SQL Server

For a planned migration where downtime is acceptable, you can create a migration task without incremental data migration. If you do not include the incremental data migration phase when you create the migration task, you must suspend all write requests to the source database during migration.

Prerequisites

• The version of the user-created SQL Server database is 2005, 2008, 2008 R2, 2012, 2014, 2016, or 2017.



- SQL Server clusters and SQL Server Always On high availability group are not supported.
- To perform full migration between databases of different versions, make sure that the two 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 user-created SQL Server database.

Background information

• 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 resumes a failed data migration 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 cannot migrate data of the sql_variant type.
- DTS does not support the schema migration for 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, and aggregate functions.

Migration types

• Schema migration

DTS migrates the schemas of objects to the destination instance. DTS can migrate the schemas of the following objects: tables, views, triggers, synonyms, SQL stored procedures, SQL functions, plan guides, user-defined types, rules, and defaults.

 Full data migration
 DTS migrates full data of the source objects from the user-created SQL Server database to the ApsaraDB RDS for SQL Server instance.

Billing

Migration type	Instance configurations	Internet traffic
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
User-created SQL Server database	SELECT privilege	SELECT privilege
ApsaraDB RDS for SQL Server database	Read and write permissions	Read and write permissions

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

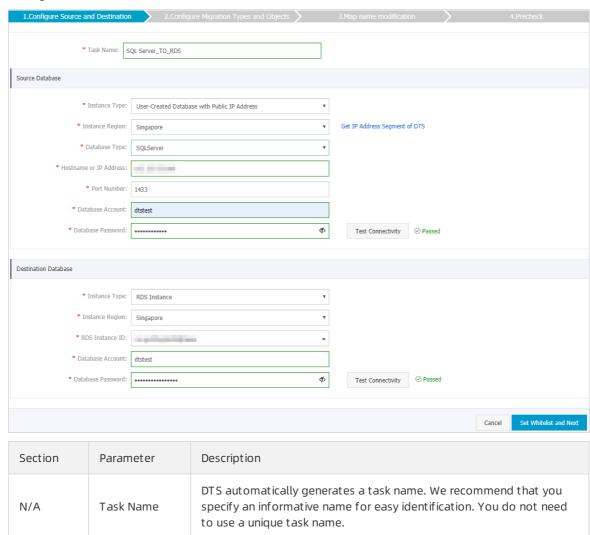
Full data migration

To avoid 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. Migrate full data.
- 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.



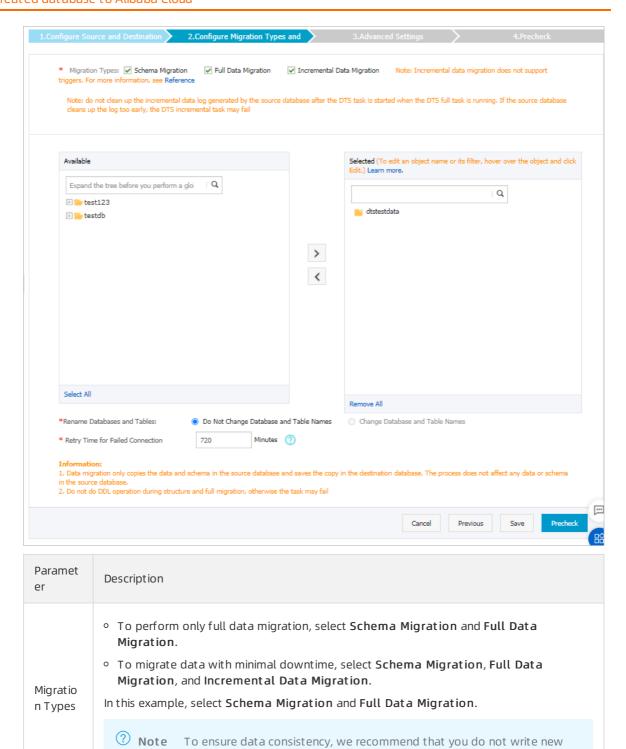
Section	Parameter	Description		
	Instance Type	Select an instance type based on the location where the database is deployed. In this topic, a User-Created Database with Public IP Address is used as an example. Note If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see Preparation overview.		
	Instance Region	If Instance Type is set to User-Created Database with Public IP Address, you do not need to specify the Instance Region. Note If you have configured a whitelist for the user-created SQL Server database, you must add the CIDR blocks of DTS servers to the whitelist. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.		
	Database Type Hostname or	Select SQL Server. Enter the endpoint that is used to connect to the user-created SQL		
Source Database Port Number Database Account Database Password	Port Number	Enter the service port number of the user-created SQL Server database. The default port number is 1433. Note In this example, the service port of the user-created MongoDB database must be open to the public network.		
	20100000	Enter the account that is used to log on to the user-created SQL Server database. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
	Phote After you specify the information about the selfmanaged 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.			

Section	Parameter	Description		
	Instance Type	Select RDS Instance.		
	Instance Region	Select the region in which the destination RDS instance resides.		
	RDS Instance ID	Select the ID of the destination RDS instance.		
Destinatio n	Dat abase Account	Enter the account that is used to log on to the destination RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.		
Database	Dat abase Password	Enter the password of the destination database account.		
		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 \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

? Note The CIDR blocks of DTS servers are automatically added 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 objects to be migrated.



data into the user-created SQL Server database during full data migration.

Paramet er	Description
	Select objects from the Available section and click the joicon to move the objects to the Selected section.
Migratio n objects	 Note The selected objects can be databases, tables, or columns. By default, the name of an object remains unchanged after migration. You can change the names of the objects in the destination RDS instance by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature for an object, objects that depend on the object may fail to be migrated.

8.

- 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 Do not manually stop a data migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task succeeds.

12. Switch your workloads to the destination RDS instance.

What to do next

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

7.3. Source: Oracle

7.3.1. Migrate data from a user-created Oracle database to a PolarDB cluster compatible with Oracle

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

Prerequisites

- The version of the user-created Oracle database is 9i, 10g, 11g, 12c, 18c, or 19c.
- The ARCHIVELOG mode is enabled for the user-created Oracle database. 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.
- Supplemental logging, including SUPPLEMENTAL_LOG_DATA_PK and SUPPLEMENTAL_LOG_DATA_UI, is enabled for the user-created Oracle database. For more information, see Supplemental Logging.
- The tables to be migrated from the user-created 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%.
- DTS automatically resumes a failed data migration 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 from Alibaba Cloud over the Internet. For more
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

Migration types

Migration type	Description
	DTS migrates the schemas of the required objects to the destination database. DTS supports schema migration for the following types of objects: table, view, synonym, trigger, stored procedure, function, package, and user-defined type.
Schema migration	Note However, if an object contains triggers, the data between the source and destination databases will become inconsistent.

Migration type	Description		
	DTS migrates historical data of the required objects from the user-created Oracle database to the destination PolarDB cluster.		
Full data migration	Note 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.		
Incremental data migration	DTS retrieves redo log files from the user-created Oracle database. Then, DTS synchronizes incremental data from the user-created Oracle database to the destination PolarDB cluster. Incremental data migration allows you to ensure service continuity when you migrate data from the user-created Oracle database to the destination PolarDB cluster.		

SQL operations that can be synchronized during incremental data migration

- INSERT, UPDATE, and DELETE operations
- CREATE TABLE operations
 - Note The CREATE TABLE operations to create partition tables or tables that contain functions cannot be synchronized.
- ALTER TABLE, DROP TABLE, RENAME TABLE, CREATE INDEX, and ADD INDEX operations

Preparations

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
User-created Oracle database	The owner permission on schemas	The owner permission on schemas	The database administrator (DBA) permission
PolarDB cluster	The owner permission on schemas	The owner permission on schemas	The owner permission on schemas

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

- CREATE USER and GRANT for a user-created Oracle database
- Create an account for a PolarDB cluster

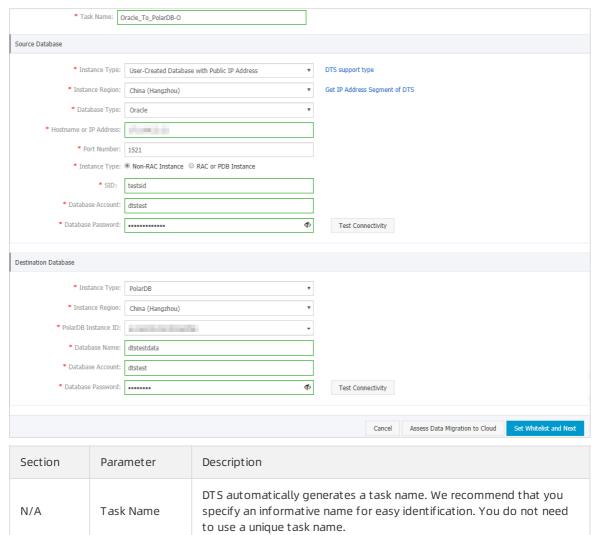
Notice If you want to migrate incremental data from an Oracle database but the DBA permission cannot be granted to the database account, you can grant fine-grained permissions to the account. The following sample statements show you how to grant specific permissions to an Oracle database account.

Grant Specific Permissions (Click to Show Details))

Sample Statements (Click to Show Details))

Procedure

- 1.
- 2.
- 3.
- 4.
- 5. Configure the source and destination databases for the data migration task.



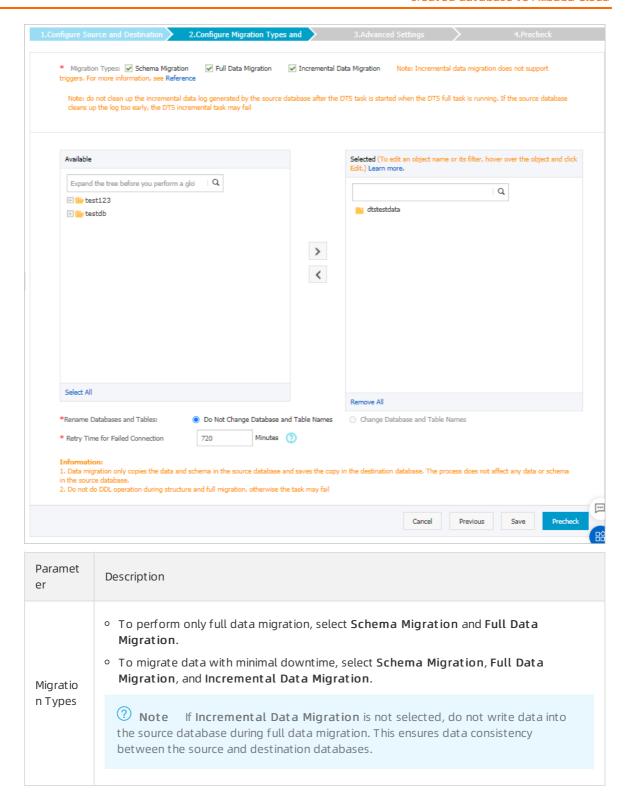
Section	Parameter	Description		
	Instance Type	Select an instance type based on where the source database is deployed. The procedure in this topic uses User-Created Database with Public IP Address as an example.		
		? Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview .		
		If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.		
	Instance Region	Note If a whitelist is configured for the user-created Oracle database, you must manually add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle 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	Enter the endpoint that is used to connect to the user-created Oracle database.		
Source Dat abase	Port Number	Enter the service port number of the user-created Oracle database. The default port number is 1521. Note In this example, The service port of the user-created Oracle database is accessible over the Internet.		
	Instance Type	 Non-RAC Instance: If you select this option, you must specify the SID. RAC Instance: If you select this option, you must specify the Service Name. 		
	Dat abase Account	Enter the account of the user-created Oracle database. For more information about the permissions that are required for the account, see Preparations.		
	Dat abase Password	Onte After you specify the information about the selfmanaged 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.		

Section	Parameter	Description	
Destinatio n Database	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.	
	Dat abase Account	Enter the database account of the destination PolarDB cluster. For more information about the permissions that are required for the account, see Preparations.	
	Dat abase Password	Enter the password for the destination database account.	
		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.

Note The CIDR blocks of DTS servers are automatically added 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 objects to be migrated.



Select objects from the Available section and click the joicon to move the objects to the Selected section.	Paramet er	
 Note You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains the same as that in the source database. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated. 	Objects	

8.

9.

10.

11.

12.

7.3.2. Migrate data from a user-created Oracle database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a user-created 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 user-created Oracle database, you can select all of the supported migration types to ensure service continuity.

Prerequisites

- The version of the user-created 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 user-created Oracle database. For more information, see Supplemental Logging.
- The ARCHIVELOG mode is enabled for the user-created Oracle database. 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 user-created Oracle database is accessible over the Internet.
- The available storage space of the destination ApsaraDB RDS for MySQL instance is larger than the total size of the data in the user-created 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.
- Table names in the ApsaraDB RDS for MySQL instance are case-insensitive. If a table name in the user-created Oracle database contains uppercase letters, ApsaraDB RDS 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 schema migration, the following message is returned: "The object already exists". To avoid name conflicts in the destination database, you can change the names of 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 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 about how to create a database and the database naming conventions, see Create a database.

Billing

Migration type	Instance configurations	Internet traffic
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.
Incremental data migration	Charged. For more 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:
 - Schema migration of nested tables is not supported. Clustered tables and index-organized tables (IOTs) are converted into common tables in the destination database.
 - Schema migration of function-based indexes, domain indexes, bitmap indexes, and reverse indexes is not supported.
- Full dat a migration
 - DTS migrates historical data of the required objects from the user-created Oracle database to the destination ApsaraDB RDS for MySQL instance.
- Increment al data migration
 DTS retrieves redo log files from the user-created Oracle database. Then, DTS synchronizes increment al data from the user-created Oracle 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 the user-created Oracle database to the destination database.

SQL operations that can be synchronized during incremental data migration

- INSERT, DELETE, and UPDATE
- CREATE TABLE
 - **?** Note The CREATE TABLE operations to create tables that contain functions cannot be synchronized.
- ALTER 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.

Preparations

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
User-created Oracle database	The owner permission on schemas	The owner permission on schemas	The database administrator (DBA) permission
ApsaraDB RDS for MySQL instance	The write permission on the destination database	The write permission on the destination database	The write permission on the destination database

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

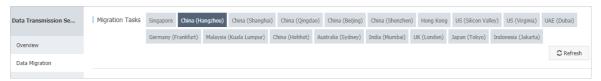
- CREATE USER and GRANT for a user-created Oracle database
- Create an account for an RDS for MySQL instance and Change the permissions of an account for an RDS for MySQL instance

Notice If you want 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.

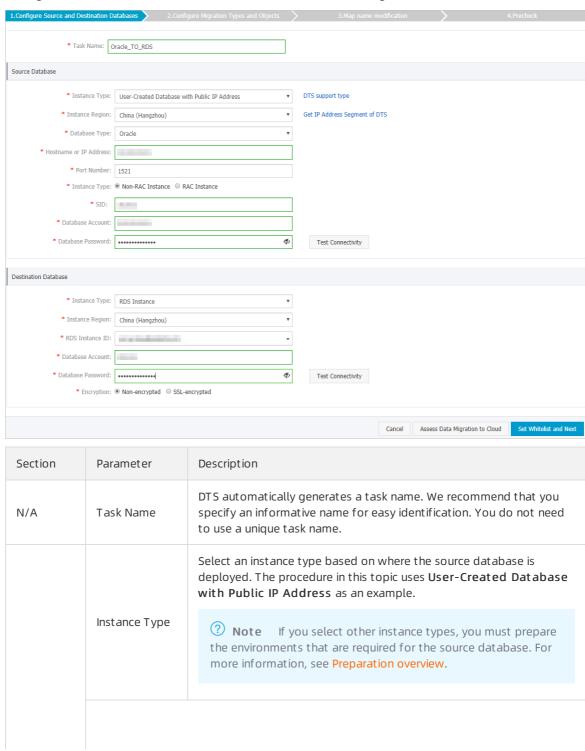
Procedure

- 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 RDS instance resides.



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



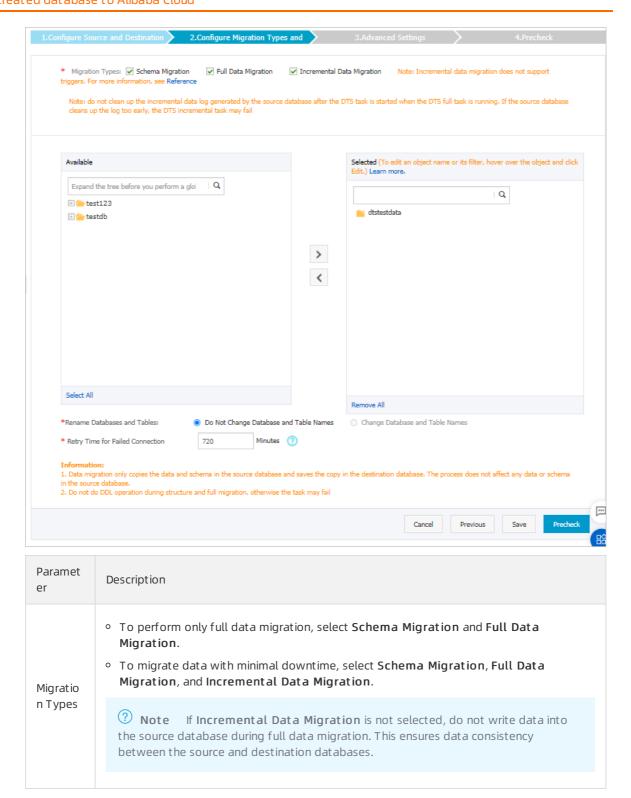
Section	Parameter	Description	
		If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region .	
	Instance Region	Note If a whitelist is configured for the user-created Oracle database, you must manually add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.	
Source	Database Type	Select Oracle.	
Database	Hostname or IP Address	Enter the IP address that is used to connect to the user-created Oracle database. In this example, enter the public IP address.	
	Port Number	Enter the service port number of the user-created Oracle database. The default port number is 1521.	
	Instance Type	 Non-RAC Instance: If you select this option, you must specify the SID. RAC Instance: If you select this option, you must specify the Service Name. 	
	Dat abase Account	Enter the account of the user-created Oracle database. For more information about the permissions that are required for the account, see Preparations.	
	Dat abase Password	Phote 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 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 more information about the permissions that are required for the account, see Preparations.	

Section Destinatio	Parameter	Description
n Dat abase		Enter the password for the destination 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 The CIDR blocks of DTS servers are automatically added 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 objects to be migrated.



113

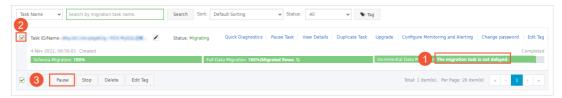
Paramet er	Description
	Select objects from the Available section and click the icon to move the objects to the Selected section.
Objects	 Note You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination RDS instance, the name of the object remains the same as that in the user-created Oracle database. You can change the names of the objects that are migrated to the destination RDS instance by using the object name mapping feature. For more information about how to use this feature, see Object name mapping.

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



- 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, click the icon next to each failed item to view details. Troubleshoot the issues based on the causes 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 and select the Data Transmission Service (Pay-As-You-Go) Service Terms.
- 11. Click Buy and Start to start the migration task.
 - Full data migration
 Do not manually stop a task during full data migration. Otherwise, data migrated to the destination database will be incomplete. Wait until the migration task automatically stops.
 - Increment al data migration
 An increment al data migration task does not automatically stop. You must manually stop the migration task.
 - Note Select an appropriate time to manually stop the migration task. For example, you can stop the migration task during off-peak hours or before you switch your workloads to the destination instance.
 - 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. After the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



12. Switch your workloads to the destination RDS instance.

What to do next

The database accounts used for data migration have the read/write permissions. After the data migration is complete, you must delete the accounts of both the user-created Oracle database and the ApsaraDB RDS for MySQL instance to ensure database security.

More information

DTS supports reverse data transmission when you migrate data from a user-created 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 user-created Oracle database. To do this, submit a ticket.

7.3.3. Migrate data from a user-created Oracle database to a DRDS instance

This topic describes how to migrate data from a user-created Oracle database to a Distributed Relational Database System (DRDS) instance by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration. When you migrate data from a user-created Oracle database, you can select the two migration types to ensure service continuity.

Prerequisites

- The version of the user-created 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 user-created Oracle database. For more information, see Supplemental Logging.
- The ARCHIVELOG mode is enabled for the user-created Oracle database. 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 user-created Oracle database is accessible over the Internet.
- The databases in DRDS instances are created based on ApsaraDB RDS for MySQL instances. DTS does not support databases that are created based on Apsara PolarDB for MySQL clusters.
- A DRDS instance is created based on an RDS instance. The available storage space of the destination RDS instance is larger than the total size of the data in the user-created Oracle database.

Precautions

- DTS does not support schema migration from a user-created Oracle database to a DRDS 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 automatically resumes a failed data migration 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.

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
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

Migration types

• Full dat a migration

DTS migrates historical data of the required objects from the user-created Oracle database to the database in the destination DRDS instance.

Note To ensure data consistency, do not write data into the source Oracle database during full data migration.

• Increment al dat a migration

After full data migration is complete, DTS retrieves redo log files from the user-created Oracle database. Then, DTS synchronizes incremental data from the user-created Oracle database to the database in the destination DRDS instance. Incremental data migration allows you to ensure service continuity when you migrate data from the user-created Oracle database to the database in the destination DRDS 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

 Create a database and tables in the destination DRDS instance based on the schema of the usercreated Oracle database. For more information, see Create a DRDS database and Create a DRDS table.

- **?** Note The data types of Oracle databases and DRDS instances do not have one-to-one correspondence. You must define the corresponding data types in DRDS instances. For more information, see Data type mappings between heterogeneous databases.
- 2. Log on to the source Oracle database, create an account for data collection, and grant permissions to the account.

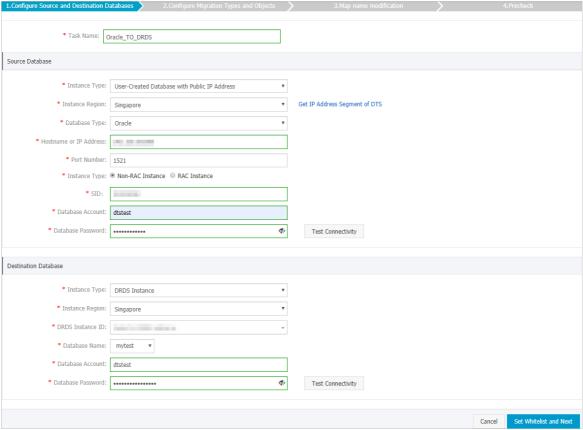
Database	Full data migration	Incremental data migration
User-created Oracle database	The owner permission on schemas	The database administrator (DBA) permission
DRDS instance	The write permission on the destination database	The write permission on the destination database

For more information about how to create and authorize an account for a user-created Oracle database, see CREATE USER and GRANT.

Notice If you want 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. 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 for the data migration task.



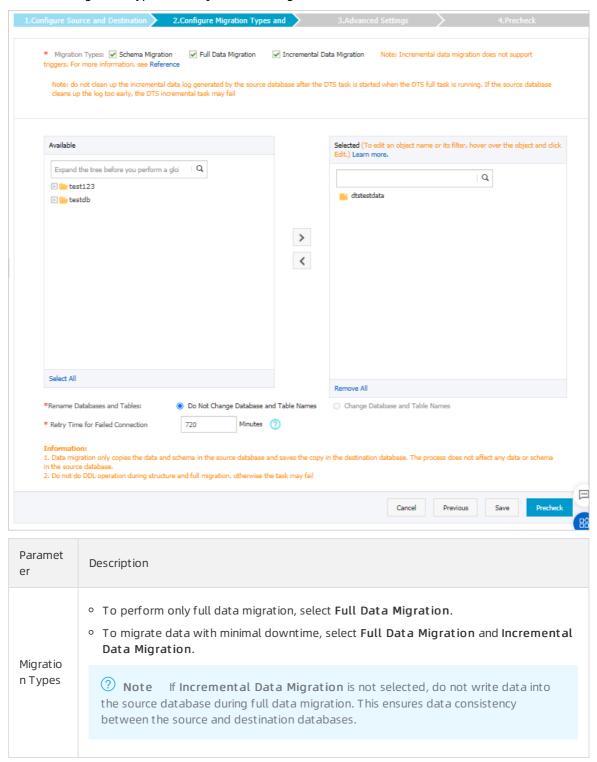
Description			
Select an instance type based on where the source database is deployed. The procedure in this topic uses User-Created Database with Public IP Address as an example. Instance Type Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview. If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region. Note If a whitelist is configured for the user-created Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.	Section	Parameter	Description
deployed. The procedure in this topic uses User-Created Database with Public IP Address as an example. 1 Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview. 1 If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region. 2 Note If a whitelist is configured for the user-created Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.	N/A	Task Name	specify an informative name for easy identification. You do not need
Instance Region Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview. If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region. Note If a whitelist is configured for the user-created Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.			deployed. The procedure in this topic uses User-Created Database
Instance Region Note If a whitelist is configured for the user-created Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.		Instance Type	the environments that are required for the source database. For
Region Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle 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 .		Oracle database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created Oracle database. You can click Get IP Address Segment of DTS next to Instance Region to	
Select Officie.		Datahase Type	Select Oracle
		Database Type	Secret G. dete.

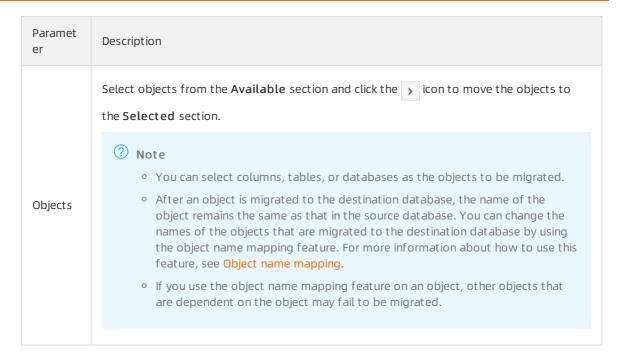
Section	Parameter	Description
Source Database	Hostname or IP Address	Enter the IP address that is used to connect to the user-created Oracle database. In this example, enter the public IP address.
	Port Number	Enter the service port number of the user-created Oracle database. The default port number is 1521 .
	Instance Type	 Non-RAC Instance: If you select this option, you must specify the SID. RAC Instance: If you select this option, you must specify the Service Name.
	Dat abase Account	Enter the account of the user-created Oracle database. For more information about the permissions that are required for the account, see Preparations.
		Enter the password for the source database account.
	Database 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	Select the region where the destination DRDS instance resides.
	DRDS Instance	Select the ID of the destination DRDS instance.
Destinatio n	Dat abase Account	Enter the database account of the destination DRDS instance. For more information about the permissions that are required for the account, see Preparations.
Database		Enter the password for the destination 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 \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination DRDS instance. This ensures that DTS servers can connect to the destination DRDS instance.

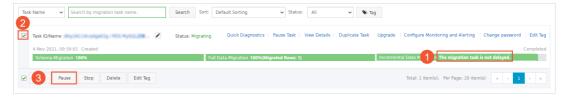
7. Select the migration types and objects to be migrated.





8.

- 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 self-managed 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 DRDS instance.

What to do next

The database accounts used for data migration have the read/write permissions. After the data migration is complete, you must delete the accounts of both the user-created Oracle database and the DRDS database to ensure database security.

7.4. Source: PostgreSQL

7.4.1. Migrate incremental data from a user-created PostgreSQL database (version 10.1 to 12) to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to migrate incremental data from a user-created 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 user-created PostgreSQL database to Alibaba Cloud, you can select all of the supported migration types to ensure service continuity.

Prerequisites

- The version of the user-created PostgreSQL database is 10.1 to 12.
- An ApsaraDB RDS for PostgreSQL instance is created. For more information, see Create an RDS for PostgreSQL instance.

? Note Ensure that the database version of the ApsaraDB RDS for PostgreSQL instance is the same as the version of the user-created PostgreSQL database.

• The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total size of the data in the user-created 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.
- To ensure that the delay time of incremental data migration is accurate, DTS adds a heart beat table named *dts_postgres_heart beat* to the source database.
- During incremental data migration, DTS creates a replication slot for the source database. The replication slot is prefixed with dts_sync. DTS automatically clears historical replication slots

every 90 minutes to prevent them from occupying disk space.

Note If a migration task is released or fails, DTS automatically clears the replication slot.
After a switchover between primary and secondary ApsaraDB RDS for PostgreSQL databases, you must log onto the secondary database to manually clear the replication slot.



• 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

•

• During incremental data migration, only data manipulation language (DML) operations, such as INSERT, UPDATE, DELETE, and UPDATE operations, can be migrated.

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
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 and authorize a database account, see the following topics:

- Self-managed PostgreSQL database: CREATE USER and GRANT
- ApsaraDB RDS for Post greSQL instance: Create an account on an ApsaraDB RDS for PPAS instance.

Data migration process

To avoid data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

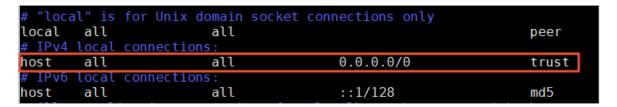
Note For more information about schema migration, full data migration, and incremental data migration, see Terms.

rates the schemas of tables, views, sequences, functions, user-defined les, domains, operations, and aggregates to the destination database.
ote Functions that are written in the C programming language t be migrated.
rates historical data of the required objects to the destination database.
rates the schemas of triggers and foreign keys to the destination e.
rates incremental data of the required objects to the destination e.
ote Incremental data migration does not support the bit data type.

Preparations

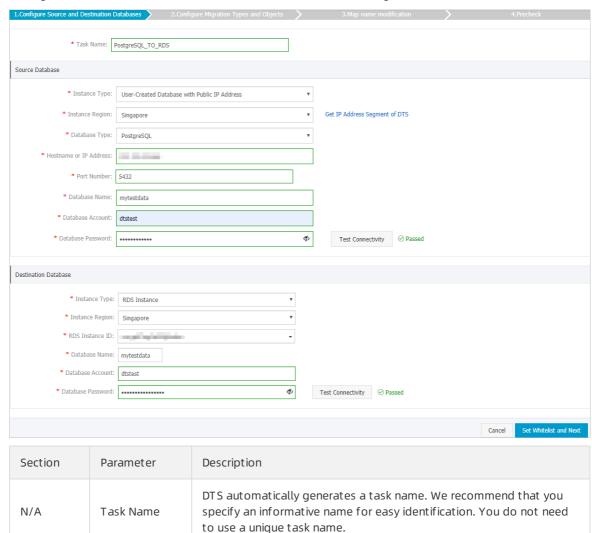
- 1. Log on to the server to which the user-created PostgreSQL database belongs.
- 2. Set the value of the wal_level parameter in the postgresql.conf configuration file to logical

- Note Skip this step if you do not need to perform incremental data migration.
- 3. Add the CIDR blocks of DTS servers to the *pg_hba.conf* configuration file of the user-created PostgreSQL database. You only need to add the CIDR blocks of DTS servers that are located 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 about how to configure the pg_hba.conf file, 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
 .0/0



Procedure

- 1.
- 2.
- 3.
- 4.
- 5. Configure the source and destination databases for the data migration task.



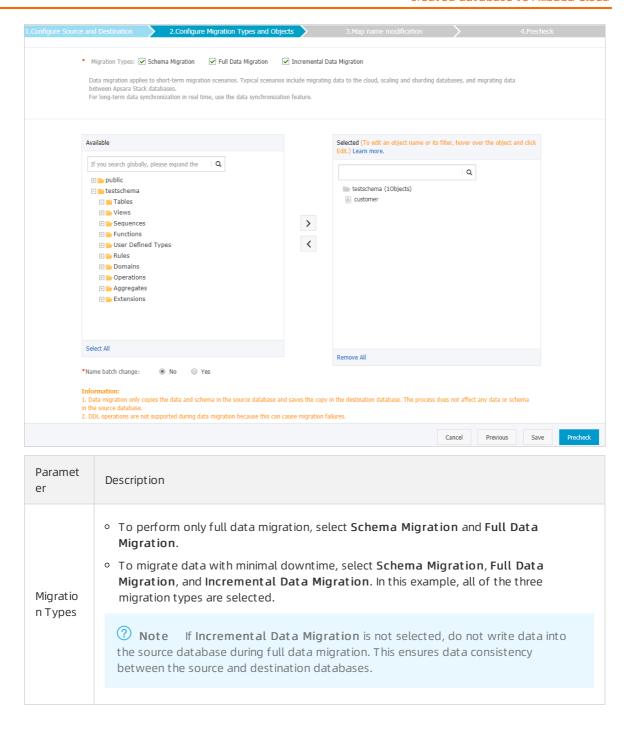
Section	Parameter	Description	
	Instance Type	Select an instance type based on where the source database is deployed. The procedure in this topic uses User-Created Database with Public IP Address as an example. Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview.	
	Instance	If the instance time is set to User Created Database with Dublic	
	Instance Region	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region .	
	Database Type	Select PostgreSQL.	
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created PostgreSQL database. In this example, enter the public IP address.	
	Port Number	Enter the service port number of the user-created PostgreSQL database. The port is accessible over the Internet.	
Source Dat abase	Dat abase Name	Enter the name of the user-created PostgreSQL database.	
	Dat abase Account	Enter the account of the user-created PostgreSQL database.	
		Enter the password for the source 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 RDS Instance.	
	Instance Region	Select the region where the destination RDS instance resides.	
	RDS Instance ID	Select the ID of the destination RDS instance.	

Parameter	Description	
	Enter the database name of the destination RDS instance. The name can be different from the name of the user-created PostgreSQL database.	
Database Name	Note Before you configure the data migration task, create a database in the ApsaraDB RDS for PostgreSQL instance. For more information, see Create a database.	
Dat abase Account	Enter the database account of the destination RDS instance.	
Dat abase Password	Enter the password for the destination database account.	
	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.	
	Database Name Database Account	

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

Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination ApsaraDB RDS for PostgreSQL instance. This ensures that DTS servers can connect to the destination ApsaraDB RDS for PostgreSQL instance.

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



Select objects from the Available section and click the icon to move the objects to the Selected section. Note You can select columns, tables, or Schemas as the objects to be migrated. If you select a Schema as an object to be migrated in a data migration task, new tables that are created after the task is started can also be migrated. However, you must run the following statement on the new tables before they can be migrated: ALTER TABLE schema.table REPLICA IDENTITY FULL; You must replace the schema and table in the preceding sample statement with the actual schema name and table name. After an object is migrated to the destination RDS instance, the name of the
 You can select columns, tables, or Schemas as the objects to be migrated. If you select a Schema as an object to be migrated in a data migration task, new tables that are created after the task is started can also be migrated. However, you must run the following statement on the new tables before they can be migrated: ALTER TABLE schema.table REPLICA IDENTITY FULL; You must replace the schema and table in the preceding sample statement with the actual schema name and table name.
 object remains the same as that in the user-created PostgreSQL database. You can change the names of the objects that are migrated to the destination RDS instance by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.

8.

9.

10.

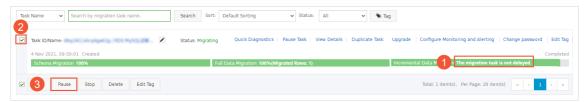
11. Click **Buy and Start** to start the migration task.

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 dat a 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.



7.4.2. Migrate incremental data from a user-created PostgreSQL database (version 9.4 to 10.0) to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to migrate incremental data from a user-created 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 user-created PostgreSQL database to Alibaba Cloud, you can select all of the supported migration types to ensure service continuity.

This topic uses **User-Created Database with Public IP Address** as an example to describe how to configure an incremental data migration task. For information about how to perform only full data migration, see Migrate full data from a self-managed PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance.

Prerequisites

- The version of the user-created PostgreSQL database is 9.4, 9.5, 9.6, 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 user-created PostgreSQL database.
- The service port of the user-created 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 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.
- DTS automatically resumes a failed data migration 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.

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
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
User-created PostgreSQL database	The usage permission on pg_catalog	The SELECT permission on the objects to be migrated	The superuser permission
ApsaraDB RDS for PostgreSQL instance	The create and usage permissions on the objects to be migrated	The owner permission on schemas	The owner permission on schemas

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

- CREATE USER and GRANT for a user-created PostgreSQL database
- Create an account for an ApsaraDB RDS for PostgreSQL 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	Note 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.	

Data migration process	Description
3. Schema migration	DTS migrates the schemas of triggers and foreign keys to the destination database.
4. Incremental data migration	DTS migrates incremental data of the required objects to the destination database. Incremental data migration ensures service continuity of self-managed applications. Note During incremental data migration, DTS migrates only data manipulation language (DML) operations. DML operations include INSERT, DELETE, and UPDATE.
	 Incremental data migration does not support the BIT data type.

Preparation 1: Install the logical flow replication plug-in and modify system settings

Before you configure an incremental data migration task, you must install the logical flow replication plug-in provided by DTS in the user-created PostgreSQL database.

- **Note** In the following procedure, PostgreSQL 9.6 is used as an example.
- 1. On the server of the user-created PostgreSQL database, run the wget command to download the logical flow replication plug-in based on the database version. Download the plug-in from one of the following links:
 - o Plug-in for PostgreSQL 9.4
 - Plug-in for PostgreSQL 9.5
 - Plug-in for PostgreSQL 9.6
 - Plug-in for PostgreSQL 10.0
- 2. Extract the plug-in from the compressed file.

```
tar xvf tar xvf ali_decoding_9.6.tar
```

3. Copy the *ali_decoding.so* file to the *lib* directory where PostgreSQL is installed.

```
cp ali decoding.so /usr/lib/postgresql/9.6/lib/
```

4. Copy the *ali_decoding.control* file to the *extension* directory where PostgreSQL is installed.

```
cp ali decoding.control /usr/share/postgresql/9.6/extension/
```

- 5. Log on to the user-created PostgreSQL database by using an account that has the superuser permission.
- 6. Set the value of the $max_replication_slots$ parameter to an integer greater than 1. In this example, set the value to 5.
 - Note For more information, see PostgreSQL official documentation.

```
ALTER SYSTEM set max_replication_slots = '5';
```

7. Set the value of the wal level parameter to logical .

```
ALTER SYSTEM SET wal_level = logical;
```

- 8. Set the value of the $max_wal_senders$ parameter to an integer greater than 1. In this example, set the value to 5.
 - ? Note The max_wal_senders parameter specifies the maximum number of concurrent tasks. We recommend that you set the value of this parameter to the same as the value of the max_replication_slots parameter.

```
ALTER SYSTEM SET max_wal_senders = '5';
```

9. Return to the shell of the server for the user-created PostgreSQL database and run the following command to restart the PostgreSQL service.

```
service postgresql restart
```

10. Log on to the user-created PostgreSQL database again and run the following command to check whether the replication slot can be created:

```
SELECT * FROM pg_create_logical_replication_slot('replication_slot_test', 'ali_decoding
');
```

The following message indicates that the logical flow replication plug-in is installed.

- 11. Add the CIDR blocks of DTS servers to the PostgreSQL configuration file (pg_hba.conf). For more information about how to edit this configuration file, see The pg_deploy.conf file. Skip this step if you have set the IP address in the pg_hba.conf file to 0.0.0.0/0.
 - Note You only need to add the CIDR blocks of 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.

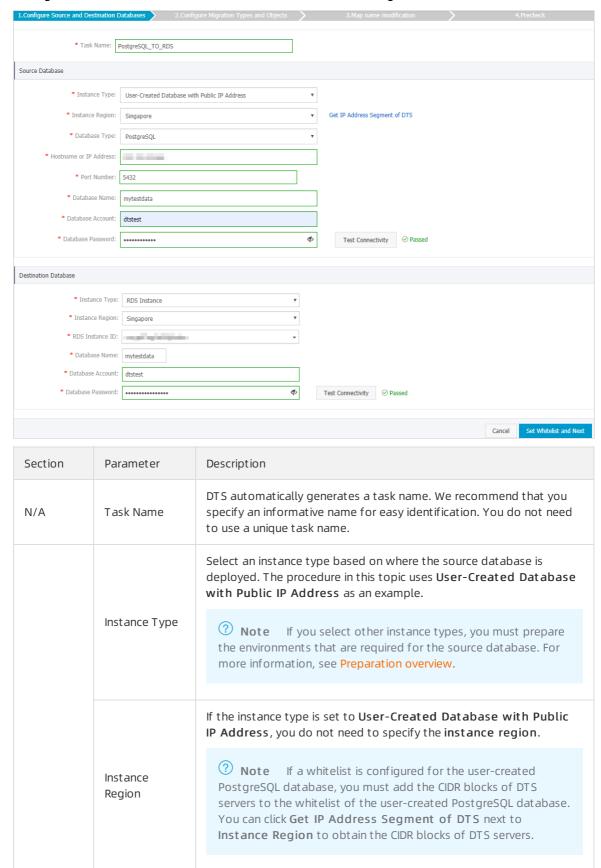
Preparation 2: Create a database and schema in the destination RDS instance

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 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.

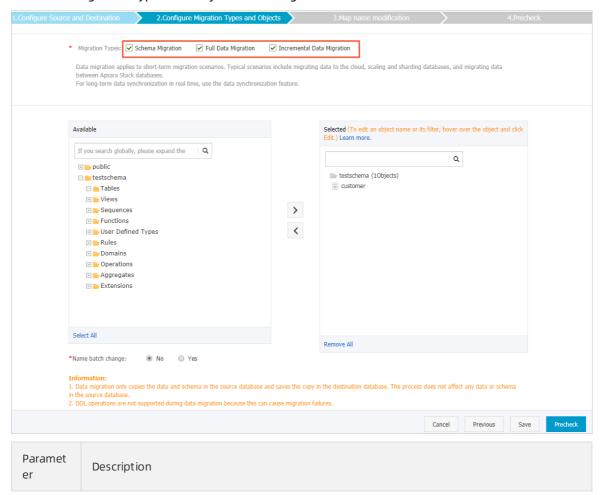
- 4. In the upper-right corner of the page, click Create Migration Task.
- 5. Configure the source and destination databases for the data migration task.



Section	Parameter	Description
Source Database	Database Type	Select PostgreSQL.
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created PostgreSQL database. In this example, enter the public IP address.
	Port Number	Enter the service port number of the user-created PostgreSQL database. The default port number is 5432 .
	Dat abase Name	Enter the name of the user-created PostgreSQL database.
	Dat abase Account	Enter the account of the user-created PostgreSQL database.
		Enter the password for the source 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 RDS Instance.
Instance Region		Select the region where the destination RDS instance resides.
	RDS Instance ID	Select the ID of the destination RDS instance.
		Enter the database name of the destination RDS instance. The name can be different from the name of the user-created PostgreSQL database.
	Dat abase Name	Note Before you configure the data migration task, create a database and schema in the destination RDS instance. For more information, see Migrate incremental data from a self-managed PostgreSQL database (in PostgreSQL 10.0 or an earlier version) to an ApsaraDB RDS for PostgreSQL instance.
Destinatio n Database		
	Database	Enter the database account of the destination RDS instance.

Section	Parameter	Description	
		Enter the password for the destination 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.
 - Note The CIDR blocks of DTS servers are automatically added 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 objects to be migrated.

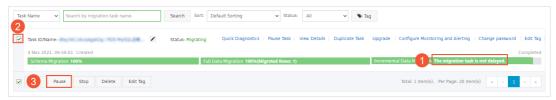


Paramet er	Description
Migratio n Types	 To perform only full data migration, select Schema Migration and Full Data Migration. To migrate data with minimal downtime, select Schema Migration, Full Data Migration, and Incremental Data Migration. In this example, all of the three migration types are selected. Note If Incremental Data Migration is not selected, do not write data into the source database during full data migration. This ensures data consistency between the source and destination databases.
Objects	Select objects from the Available section and click the icon to move the objects to the Selected section. Note You can select columns, tables, or Schemas as the objects to be migrated. You can select a Schema as an object to be migrated, and then create a new table in the source PostgreSQL database during data migration. In this case, you must run the following statement on the new table: ALTER TABLE schema.table REPLICA IDENTITY FULL; Replace the schema and table in the preceding sample statement with the actual schema name and table name. After an object is migrated to the destination RDS instance, the name of the object remains the same as that in the user-created PostgreSQL database. You
	can change the names of the objects that are migrated to the destination RDS instance by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. o If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.

8.

- 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 RDS instance.

What to do next

The database accounts used for data migration have the read/write permissions. After the data migration is complete, you must delete the accounts of both the user-created PostgreSQL database and the ApsaraDB RDS for PostgreSQL instance to ensure database security.

7.4.3. Migrate full data from a user-created PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to migrate full data from a user-created 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 user-created PostgreSQL database, you can select schema migration and full data migration.

Background

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 into the user-created PostgreSQL database during full data migration. 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.

Prerequisites

- The version of the user-created PostgreSQL database is 9.2, 9.3, 9.4, 9.5, 9.6, 10.x, 11, or 12.
- The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total size of the data in the user-created PostgreSQL database.
- The service port of the user-created 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 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.
- DTS automatically resumes a failed data migration 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 instance. DTS supports schema migration for the following types of objects: table, trigger, view, sequences function, user-defined type, rule, domain, operation, and aggregate.
- Full data migration
 DTS migrates historical data of the required objects from the user-created PostgreSQL database to the destination ApsaraDB RDS for PostgreSQL instance.

Billing

Migration type	Instance configurations	Internet traffic
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
User-created PostgreSQL database	The usage permission on pg_catalog	The permission to perform SELECT operations on the objects to be migrated
ApsaraDB RDS for PostgreSQL instance	The create and usage permissions on the objects to be migrated	The owner permission on schemas

For more 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 avoid data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

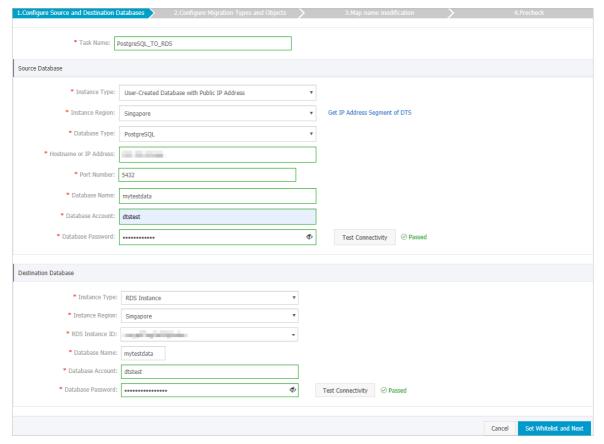
- 1. Migrate the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates.
 - **?** Note Functions that are written in the C programming language cannot be migrated.
- 2. Migrate full data.
- 3. Migrate the schemas of triggers and foreign keys.

Preparations before data migration

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 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.



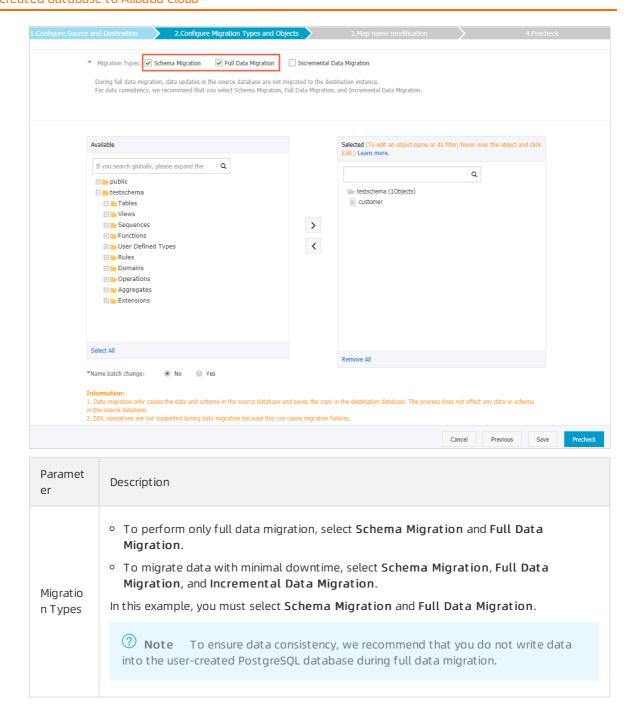
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select an instance type based on where the source database is deployed. The procedure in this topic uses User-Created Database with Public IP Address as an example. Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview.
	Instance Region	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region. ② Note If a whitelist is configured for the user-created PostgreSQL database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created PostgreSQL database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
Source	Database Type	Select PostgreSQL.
Database		Enter the endpoint that is used to connect to the user-created PostgreSQL database. In this example, enter the public IP address.
		Enter the service port number of the user-created PostgreSQL database. The default port number is 5432 .
		Enter the name of the user-created PostgreSQL database.
	Dat abase Account	Enter the account of the user-created PostgreSQL database.
	Dat abase Password	Enter the password for the source database account.
		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 RDS Instance.
	Instance Region	Select the region where the destination RDS instance resides.

Section	Parameter	Description
Destinatio n Database	RDS Instance ID	Select the ID of the destination RDS instance.
	Database Name	Enter the database name of the destination RDS instance. The name can be different from the name of the user-created PostgreSQL database.
		Note 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.
	Dat abase Password	Enter the password for the destination database account.
		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 \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

Note The CIDR blocks of DTS servers are automatically added 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 objects to be migrated.



Paramet er	Description		
Objects	Select objects from the Available section and click the icon to move the objects to the Selected section.		
	 Note You can select columns, tables, or Schemas as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains the same as that in the source database. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated. 		

8.

- 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 Do not manually stop a task during full data migration. Otherwise, data migrated to the destination database will be incomplete. Wait until the migration task automatically stops.
- 12. Switch your workloads to the destination RDS instance.

What to do next

The database accounts used for data migration have the read/write permissions. After the data migration is complete, you must delete the accounts of both the user-created PostgreSQL database and the ApsaraDB RDS for PostgreSQL instance to ensure database security.

7.5. Source: Redis

7.5.1. Migrate data from a user-created Redis database to an ApsaraDB for Redis instance

This topic describes how to migrate data from a user-created 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 user-created Redis database to Alibaba Cloud, you can select the two migration types to ensure service continuity.

Prerequisites

- The version of the user-created Redis database is 2.8, 3.0, 3.2, 4.0, or 5.0.
- The user-created Redis database uses the standalone architecture rather than the cluster architecture.

Note If the user-created Redis database uses 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 user-created Redis database.
- The available storage space of the destination ApsaraDB for Redis database is larger than the total size of the data in the user-created Redis database.

Precautions

- DTS uses resources of the source and destination databases during full data migration. This may increase the load of the database server. If the data volume is large or the specification is low, database services may become unavailable. Before you migrate data, evaluate the performance of the source and destination databases. We recommend that you migrate data during off-peak hours.
- If the data eviction policy (maxmemory-policy) of the destination database is not set to noevicti on , the data between the source and destination databases may become inconsistent. For more information about the data eviction policy, see How does ApsaraDB for Redis evict data by default?
- If you use 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 calling the PSYNC or SYNC command to transmit data of the LIST type, DTS does not perform the flush operation on the existing data. Therefore, duplicate data may exist in the destination database.
- DTS automatically resumes a failed data migration 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.

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
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 user-created Redis database to the destination ApsaraDB for Redis instance.

Note To ensure data consistency, do not write data into the user-created Redis database during full data migration.

• Incremental data migration

After full data migration is complete, DTS synchronizes incremental data from the user-created Redis database to the destination ApsaraDB for Redis instance. Incremental data migration allows you to ensure service continuity when you migrate data from a user-created Redis database to Alibaba Cloud.

Operations that can be synchronized during incremental data migration

- APPEND
- BIT OP, BLPOP, BRPOP, and BRPOPLPUSH
- DECR, DECRBY, and DEL
- EVAL, EVALSHA, EXEC, EXPIRE, and EXPIREAT
- FLUSHALL and FLUSHDB
- GEOADD and GETSET
- HDEL, HINCRBY, HINCRBYFLOAT, HMSET, HSET, and HSETNX
- INCR, INCRBY, and INCRBYFLOAT
- LINSERT, LPOP, LPUSH, LPUSHX, LREM, LSET, and LTRIM
- 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 incremental data migration tasks can run 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.

- Note Skip this step if you perform only full dat a migration.
- 1. Use the redis-cli program to connect to the user-created Redis database.
 - **Note** You can use the redis-cli program after you install the Redis client. For more information, see Redis community official website.

redis-cli -h <host> -p <port> -a <password>

? Note

- <host>: the endpoint that is used to connect to the user-created Redis database. You
 can use 127.0.0.1 in this example.
- <port>: the service port number of the user-created Redis database. The default port number is 6379.
- o <password>: the password for the user-created Redis database.

Example

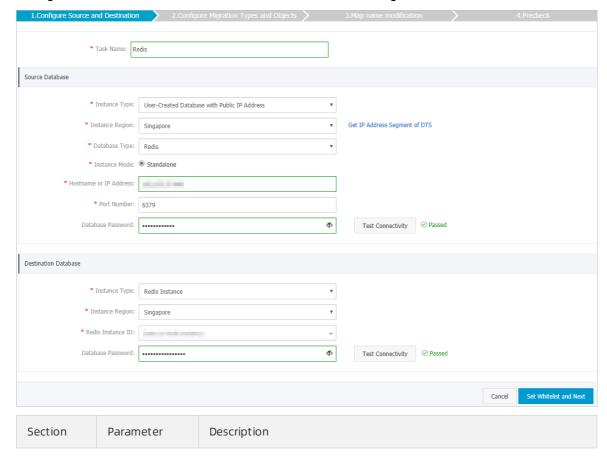
```
redis-cli -h 127.0.0.1 -p 6379 -a Test123456
```

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 for the data migration task.



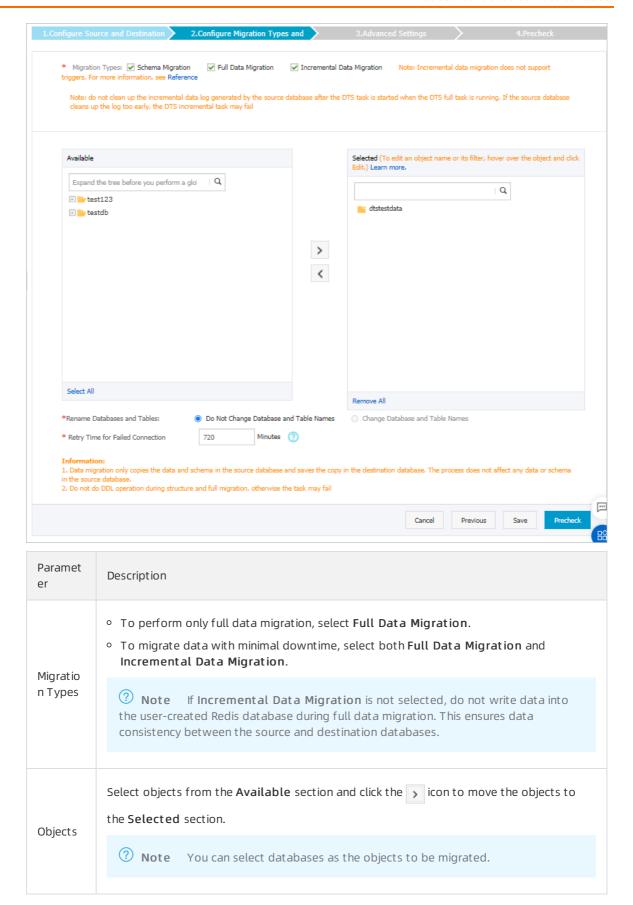
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select an instance type based on where the source database is deployed. The procedure in this topic uses User-Created Database with Public IP Address as an example. Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview.
	lnst ance Region	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region. ? Note If a whitelist is configured for the user-created Redis database, you must add the CIDR blocks of DTS servers to the whitelist of the user-created Redis database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.
Source Database	Database Type	Select Redis .
	Instance Mode	The value of this parameter is set to Standalone and cannot be changed to Cluster.
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created Redis database. In this example, enter the public IP address.
	Port Number	Enter the service port number of the user-created Redis database. The default port number is 6379.
		Enter the password for the user-created Redis database.
	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 Redis Instance .
	Instance Region	Select the region where the destination ApsaraDB for Redis instance resides.

Section	Parameter	Description
Destinatio n	natio Redis Instance ID Select the ID of the destination ApsaraDB for Redis in	
Database		Enter the database password of the destination ApsaraDB for Redis instance.
	Database 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 \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

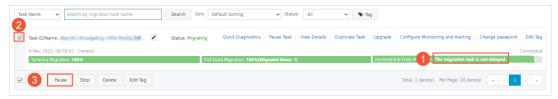
Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination ApsaraDB for Redis instance. This ensures that DTS servers can connect to the destination ApsaraDB for Redis instance.

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



8.

- 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 ApsaraDB for Redis instance.

What to do next

The database accounts used for data migration have the read/write permissions. After the data migration is complete, you must delete the accounts of both the user-created Redis database and the ApsaraDB for Redis instance to ensure database security.

7.6. Source: MongoDB

7.6.1. Migrate user-created standalone MongoDB databases to Alibaba Cloud by using DTS

This topic describes how to use Data Transmission Service (DTS) to migrate data from standalone user-created MongoDB databases to Alibaba Cloud. DTS supports full data migration and incremental data migration.

To migrate all data without service interruption, you can select both full data migration and incremental data migration. You can also use the built-in commands of MongoDB to migrate user-created 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 Data migration procedure list.

Prerequisites

- The version of the user-created MongoDB database is 3.0, 3.2, or 3.4.
- The storage space of the ApsaraDB for MongoDB instance is larger than the size of the user-created MongoDB database.

Precautions

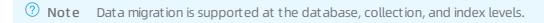
- To avoid service disruptions, we recommend that you migrate data during off-peak hours.
- If the source user-created MongoDB database and the destination ApsaraDB for MongoDB instance run different MongoDB versions or storage engines, ensure that your applications can run on both databases. For more information about MongoDB versions and storage engines that are supported by ApsaraDB for MongoDB, see MongoDB versions and storage engines.
- To migrate incremental data from a standalone user-created MongoDB database, you must enable the oplog feature for the database. For more information, see Preparations for incremental data migration.

Billing

Migration type	Instance configurations	Internet traffic
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 historical data in the source MongoDB database is migrated to the destination MongoDB database.



• Incremental data migration: After full data migration, incremental data is synchronized to the destination MongoDB database.



- The create and delete operations for databases, collections, and indexes can be synchronized.
- The create, delete, and update operations for documents can also be synchronized.

Permissions required for database accounts

Database	Full data migration	Incremental data migration
Source user-created MongoDB database	Read permissions on the source database	Read permissions on the source database, admin database, and local database

Database	Full data migration	Incremental data migration
Destination ApsaraDB for MongoDB instance	Read/write permissions on the destination database	Read/write permissions on the destination database

For more information about how to create and authorize a database account:

- Create User for MongoDB for a user-created MongoDB database
- Manage MongoDB users though DMS for an ApsaraDB for MongoDB instance

Preparations for incremental data migration

Before you use DTS to migrate incremental data, enable the oplog feature for the source database. If you only perform full data migration, skip the following steps.

Note This operation restarts the MongoDB database. We recommend that you perform this operation during off-peak hours.

- 1. Use Mongo Shell to connect to the user-created Mongo DB database.
- 2. Run the following commands to shut down the MongoDB database:

```
use admin
db.shutdownServer()
```

3. Run the following command to start the MongoDB database in the background as a replica set:

```
mongod --port 27017 --dbpath /var/lib/mongodb --logpath /var/log/mongodb/mongod.log --r
eplSet rs0 --bind_ip 0.0.0.0 --auth --fork
```

? Note

- In this command, /var/lib/mongodb is the database path, and /var/log/mongodb/mongod.log is the log file path. Specify the paths based on business needs.
- o This command uses 0.0.0.0 as the associated IP address of the MongoDB database. This allows you to access the database by using all IP addresses. After the migration is complete, run the kill command to end the process, and start the MongoDB database by using the original configuration file.
- This command enables authentication. You can only access the database after you pass the authentication.
- 4. Use Mongo Shell to reconnect to the user-created MongoDB database.
- 5. Run the following commands to initialize the replica set:

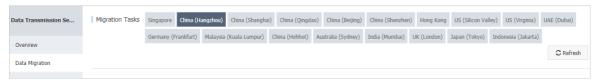
```
use admin
rs.initiate()
```

6. The role of the current node changes to primary.

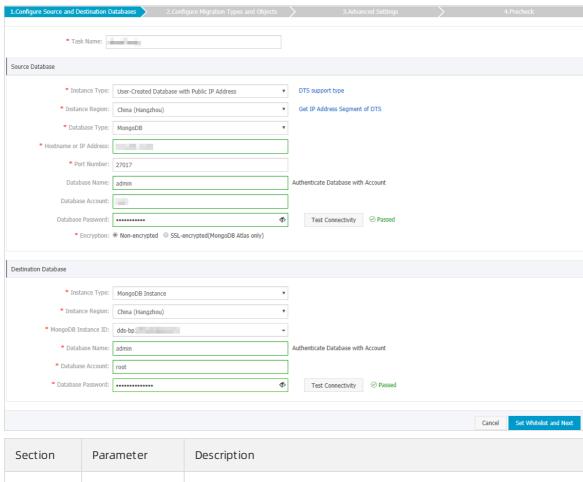
```
Note You can run the rs.printReplicationInfo() command to view the status of oplog.
```

Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the **Migration Tasks** section, select the region in which the ApsaraDB for MongoDB instance resides.



- 4. In the upper-right corner, 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 for easy identification. You do not need to use a unique task name.

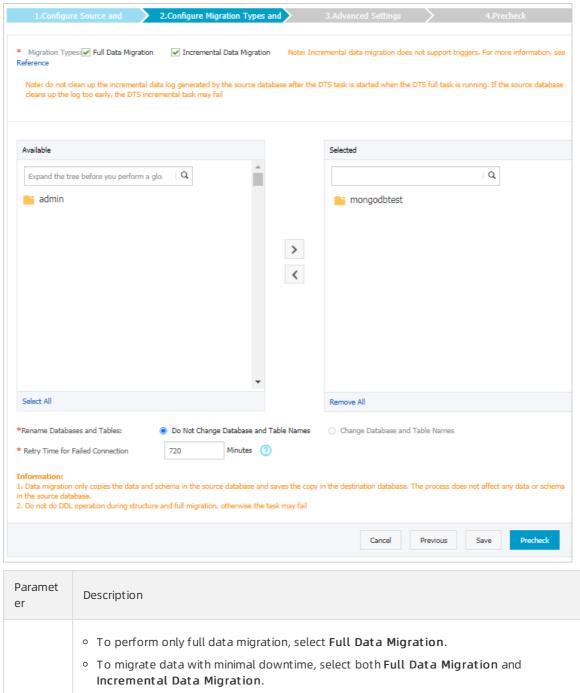
Section	Parameter	Description
	Instance Type	Select an instance type based on the location where the database is deployed. In this topic, a User-Created Database with Public IP Address is used as an example. Onote If you select other instance types, you must deploy
		the network environment for the self-managed database. For more information, see Preparation overview.
		If Instance Type is set to User-Created Database with Public IP Address , you do not need to specify the Instance Region parameter.
	Instance Region	Note If you have configured a whitelist for the user-created MongoDB database, you must add the CIDR blocks of DTS servers to the whitelist. 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.
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created MongoDB database. In this example, enter the public IP address.
Source Dat abase	Port Number	Enter the service port of the user-created MongoDB database. Note In this example, the service port of the user-created MongoDB database must be open to the public network.
	Dat abase Name	Enter the name of the authentication database to which the database account belongs.
	Database Account	Enter the username of the database account that you use to manage the source database.
	Dat abase Password	Onte After you specify the source database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the Failed message is displayed, you can click Check next to the Failed message to modify the information as prompted.
		click Test Connectivity next to Database Password to whether the information is correct. If the information is cothe Passed message is displayed. If the Failed message displayed, you can click Check next to the Failed message

Section	Parameter	Description	
		Select Non-encrypted.	
	Connection Method	Note The SSL-encrypted option is available only when you migrate MongoDB Atlas.	
	Instance Type	Select MongoDB Instance.	
	Instance Region	Select the region in which the destination ApsaraDB for MongoDB instance resides.	
	MongoDB Instance ID	Select the ID of the destination ApsaraDB for MongoDB instance.	
Destinatio n Database	Dat abase Name	Enter the name of the authentication database to which the database account belongs. Onote If you want to use the root account, specify admin for the Database Name parameter.	
Database	Dat abase Account	Enter the username of the database account that you use to manage the destination database.	
	Dat abase Password	Enter the password of the database account. Note After you specify the destination database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the Failed message is displayed, click Check next to the Failed message to modify the information as prompted.	

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

Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination RDS instance. This ensures that DTS servers can connect to the destination ApsaraDB for MongoDB instance. After the migration is completed, you can remove these CIDR blocks from the whitelist. For more information, see Configure a whitelist for a replica set instance.

7. Configure migration types and objects to be migrated.



Paramet er

Description

To perform only full data migration, select Full Data Migration.

To migrate data with minimal downtime, select both Full Data Migration and Incremental Data Migration.

Note

Before migrating incremental data from a standalone user-created MongoDB database, you must enable oplog for the database. For more information, see Preparations for incremental data migration.

If the Incremental Data Migration option is not selected, do not write new data to the user-created MongoDB database when full data migration is in progress. Otherwise, data inconsistency may occur.

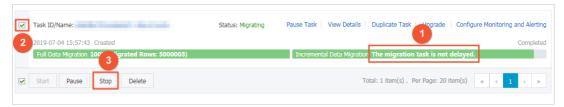
Paramet er	Description	
	 Select objects from the Available section and click the icon to move the objects to the Selected section. 	
Migratio n objects	 Note Data in the admin and local databases cannot be migrated. The config database is an internal database. We recommend that you do not migrate data in the config database. 	
	 A migration object can be a database, collection, or function. By default, the name of an object remains unchanged after migration. You can change the names of the objects in the destination RDS instance by using the object name mapping feature. For more information, see Object name mapping. 	

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



- 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.
- o If the task fails to pass the precheck, click the icon next to each failed item to view details. Troubleshoot the issues based on the causes 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 migration task.
 - Full data migration
 Do not manually stop a task during full data migration. Otherwise, the system may fail to perform a full data migration. Wait until the data migration task automatically stops.
 - Incremental data migration
 An incremental data migration task does not automatically stop. You must manually stop the migration task.
 - Note Select an appropriate time to manually stop the migration task. For example, you can stop the migration task during off-peak hours or before you switch your workloads to the destination instance.
 - 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. After the status of Incremental Data Migration changes to The migration task is not delayed, stop the migration task.



12. Switch your workloads to the ApsaraDB for MongoDB instance.

7.6.2. Migrate the replica set of a user-created MongoDB database to ApsaraDB for MongoDB by using DTS

This topic describes how to migrate the replica set of a user-created MongoDB database to ApsaraDB for MongoDB by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration.

To migrate all data without service interruption, you can select both full data migration and incremental data migration. You can also use the built-in commands of MongoDB to migrate user-created 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 or synchronization solutions, see Overview.

Prerequisites

- The version of the user-created MongoDB database is 3.0, 3.2, 3.4, 3.6, or 4.0.
- The storage capacity of the ApsaraDB for MongoDB instance is larger than the size of the user-created MongoDB database.

Precautions

- During full data migration, DTS occupies some storage space of the source and destination databases. This may increase the load of the database servers. If you migrate a large volume of data or the server specifications cannot meet your requirements, the databases may be overloaded or become unavailable. We recommend that you migrate user-created MongoDB databases during offpeak hours.
- If the source user-created MongoDB database and the destination ApsaraDB for MongoDB instance run different MongoDB versions or storage engines, ensure that your applications can run on both databases. For more information about MongoDB versions and storage engines that are supported by ApsaraDB for MongoDB, see MongoDB versions and storage engines.

Billing

Migration type	Instance configurations	Internet traffic
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, visit Pricing.

Migration types

- Full data migration: All historical data in the source MongoDB database is migrated to the destination MongoDB database.
 - Note Data migration is supported at the database, collection, and index levels.
- Incremental data migration: After full data migration, incremental data is synchronized to the destination MongoDB database.
 - ? Note
 - The create and delete operations for databases, collections, and indexes can also be synchronized.
 - The create, delete, and update operations for documents can be synchronized.

Permissions required for database accounts

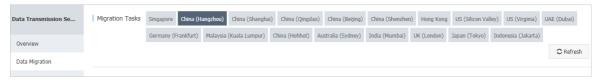
Database	Full data migration	Incremental data migration
Source user-created MongoDB database	Read permissions on the source database	Read permissions on the source database, admin database, and local database
Destination ApsaraDB for MongoDB database	Read/write permissions on the destination database	Read/write permissions on the destination database

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

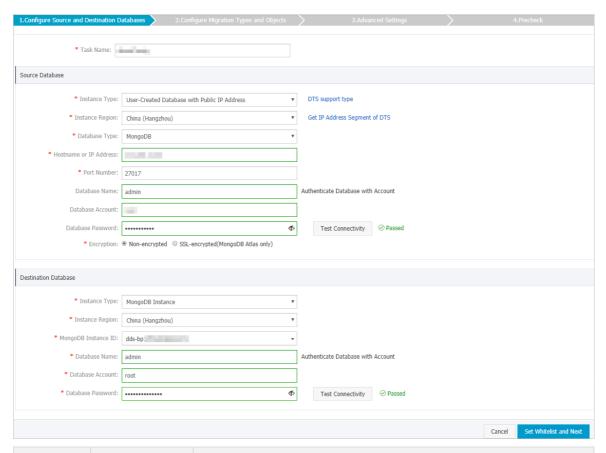
- Manage MongoDB users through DMS for an ApsaraDB for MongoDB instance
- Create User in MongoDB for a user-created MongoDB database

Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Migration**.
- 3. In the **Migration Tasks** section, select the region in which the ApsaraDB for MongoDB instance resides.



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



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 an instance type based on the location where the database is deployed. In this topic, a User-Created Database with Public IP Address is used as an example.	
		Note If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see Preparation overview.	

Section	Parameter	Description	
Source Dat abase	Instance Region	If Instance Type is set to User-Created Database with Public IP Address, you do not need to specify the Instance Region parameter. ? Note If you have configured a whitelist for the user-created MongoDB database, you must add the CIDR blocks of DTS servers to the whitelist. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.	
	Dat abase Type	Select MongoDB.	
	Hostname or IP Address	Enter the IP address of the user-created MongoDB database. In this example, enter the public IP address.	
	Port Number	Enter the service port of the user-created MongoDB database. Note In this example, the service port of the user-created MongoDB database must be open to the public network.	
	Dat abase Name	Enter the name of the authentication database to which the database account belongs.	
	Dat abase Account	Enter the account that is used to log on to the user-created MongoDB database.	
	Dat abase Password	Enter the password of the source database account. Note After you specify the source database information, click Test Connectivity next to Database Password to check whether the information is correct. If the source database configuration is correct, the Passed message is displayed. If the Test Failed message is displayed, click Check next to the Failed message to modify the information as prompted.	
	Connection Method	Select Non-encrypted. Note The SSL-encrypted option is available only when you migrate MongoDB Atlas.	
	Instance Type	Select MongoDB Instance.	
	Instance Region	Select the region in which the ApsaraDB for MongoDB instance resides.	

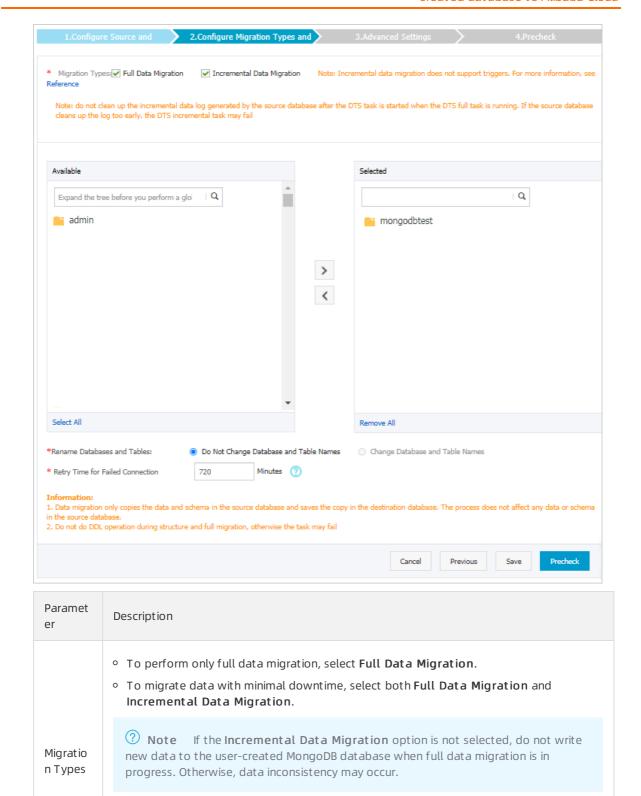
Section	Parameter	Description	
	MongoDB Instance ID	Select the ID of the ApsaraDB for MongoDB database.	
	Database Name	Enter the name of the authentication database to which the database account belongs.	
Destinatio n Database		Note If you want to use the root account, specify admin for the Database Name parameter.	
	Dat abase Account	Enter the account that is used to log on to the ApsaraDB for MongoDB instance.	
	Database Password	Enter the password of the destination database account.	
		Note After you specify the destination database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed, and you must click Check next to the Failed message to modify the information as prompted.	

6. In the lower-right corner of the page, click \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

? Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination RDS instance. This ensures that DTS servers can connect to the destination ApsaraDB for MongoDB instance. After the migration is completed, you can remove these CIDR blocks from the whitelist. For more information, see Configure a whitelist for an ApsaraDB for MongoDB instance.

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

163



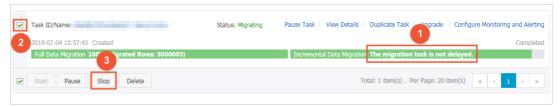
Paramet er	Description
Migratio n objects	 Select objects from the Available section and click the icon to move the objects to the Selected section. Note Data in the admin and local databases cannot be migrated. The config database is an internal database. We recommend that you do not migrate data in the config database.
	 A migration object can be a database, collection, or function. By default, the name of an object remains unchanged after migration. You can change the names of the objects in the destination RDS instance by using the object name mapping feature. For more information, see Object name mapping.

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



- 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.
- o If the task fails to pass the precheck, click the icon for each failed check item to view details. Troubleshoot the issues based on the causes 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 and select the Data Transmission Service (Pay-As-You-Go) Service Terms.
- 11. Click Buy and Start to start the migration task.
 - Full data migration
 Do not manually stop a task during full data migration. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task succeeds.
 - Incremental data migration
 An incremental data migration task does not automatically stop. You must manually stop the migration task.
 - Note Select an appropriate time to manually stop the migration task. For example, you can stop the migration task during off-peak hours or before you switch your workloads to the destination instance.
 - 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. After the status of Incremental Data Migration changes to The migration task is not delayed, stop the migration task.



12. Switch your workloads to the ApsaraDB for MongoDB instance.

References

How do I connect to the replica set instance of ApsaraDB for MongoDB?

7.6.3. Migrate the user-created sharded MongoDB database to ApsaraDB for MongoDB by using DTS

This topic describes how to migrate shards of a user-created MongoDB database to ApsaraDB for MongoDB by using Data Transmission Service (DTS). DTS allows you to migrate historical and incremental data without service disruptions.

For more information about data migration and synchronization solutions, see Overview.

Prerequisites

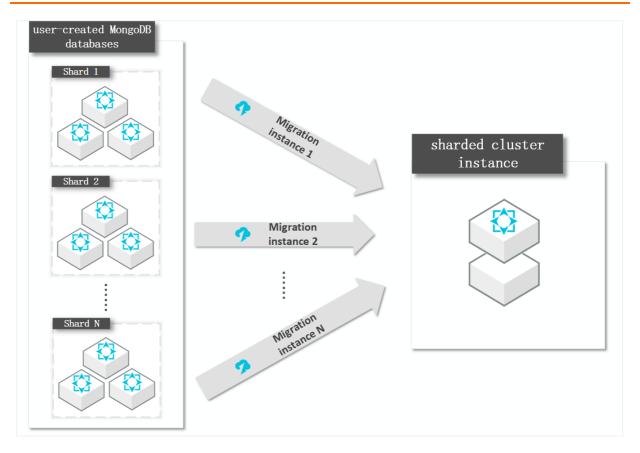
- The version of the user-created MongoDB database is 3.0, 3.2, 3.4, 3.6, or 4.0.
- Each shard in the ApsaraDB for MongoDB instance has sufficient storage space.

? Note For example, a user-created MongoDB database has three shards, and one of these shards occupies the most storage space (500 GB). In this case, the storage space of each shard in the ApsaraDB for MongoDB instance must be greater than 500 GB.

How it works

DTS migrates a user-created 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 ApsaraDB for MongoDB instance depends on the shard key that you specify. For more information, see Configure sharding to maximize the performance of shards.



Precautions

- During full data migration, DTS occupies some storage space of the source and destination databases. This may increase the load of the database servers. If you migrate a large volume of data or the server specifications cannot meet your requirements, the databases may be overloaded or become unavailable. We recommend that you migrate user-created MongoDB databases during offpeak hours.
- If the source user-created MongoDB database and the destination ApsaraDB for MongoDB instance run different MongoDB versions or storage engines, ensure that your applications can run on both databases. For more information about MongoDB versions and storage engines that are supported by ApsaraDB for MongoDB, see MongoDB versions and storage engines.

Billing

Migration type	Instance configurations	Internet traffic
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 public network. For more information, visit Pricing.

Migration types

• Full data migration: All historical data in the source MongoDB database is migrated to the destination MongoDB database.

Note Data migration is supported at the database, collection, and index levels.

• Incremental data migration: After full data migration, incremental data is synchronized to the destination MongoDB database.



- The create and delete operations for databases, collections, and indexes can also be synchronized.
- The create, delete, and update operations for documents can be synchronized.

Permissions required for database accounts

Database	Full data migration	Incremental data migration
Source user-created MongoDB database	Read permissions on the source database	Read permissions on the source database, admin database, and local database
Destination ApsaraDB for MongoDB instance	Read/write permissions on the destination database	Read/write permissions on the destination database

For more information about how to create and authorize a database account:

- Create User in MongoDB for a user-created MongoDB database.
- Manage MongoDB users through DMS for a sharded cluster instance of ApsaraDB for MongoDB.

Preparations

- 1. Disable the balancer of the user-created MongoDB database. For more information, see Manage the ApsaraDB for MongoDB balancer.
- 2. Delete the orphaned documents generated due to migration failures from the user-created MongoDB database.
 - Note If the orphaned documents are not deleted, the documents with __id conflicts may exist during migration and unwanted data may be migrated.
 - i. Download the cleanupOrphaned.js file.

- ii. In the *cleanupOrphaned.js* file, replace test with the name of the database from which you want to delete orphaned documents.
 - Note If you want to delete orphaned documents from multiple databases, repeat step ii and step iii.

- iii. Run the following command on a shard to delete the orphaned documents from all collections in the specified database:
 - Note You must repeat this step for each shard.

```
mongo --host <Shardhost> --port <Primaryport> --authenticationDatabase <database>
-u <username> -p <passowrd> cleanupOrphaned.js
```

? Note

- <Shardhost>: the IP address of the shard
- <Primaryport>: the service port of the primary node of the shard
- <database>: the name of the authentication database to which the database account belongs
- <username>: the username that you use to log on to the user-created MongoDB database
- <password>: the password that you use to log on to the user-created MongoDB database

Example:

In this example, a user-created 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

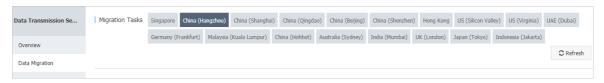
mongo --host 172.16.1.11 --port 27021 --authenticationDatabase admin -u root -p 'Te st123456' cleanupOrphaned.js

mongo --host 172.16.1.12 --port 27024 --authenticationDatabase admin -u root -p 'T est123456' cleanupOrphaned.js
```

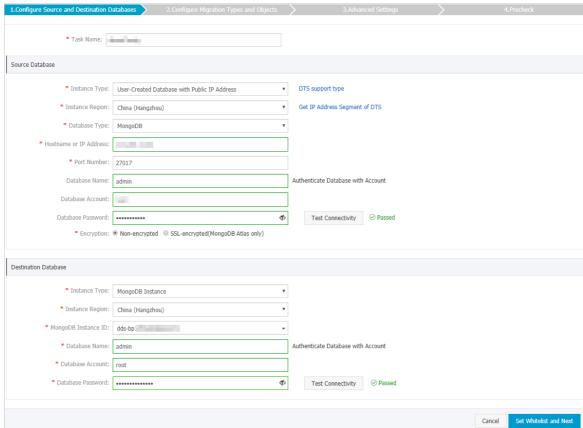
- 3. Create required databases and collections in the destination ApsaraDB for MongoDB instance, and configure data sharding for the databases and collections. For more information, see Configure sharding to maximize the performance of shards.
 - Note If you configure data sharding before you start data migration, data in the user-created MongoDB database is evenly migrated to the shards in the destination sharded cluster instance. This prevents overloading a single shard.

Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Dat a Migration**.
- 3. In the **Migration Tasks** section, select the region in which the ApsaraDB for MongoDB instance resides.



- 4. In the upper-right corner, 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 for easy identification. You do not need to use a unique task name.	
	Instance Type	Select an instance type based on the location where the database is deployed. In this topic, a User-Created Database with Public IP Address is used as an example.	
		Note If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see Preparation overview.	

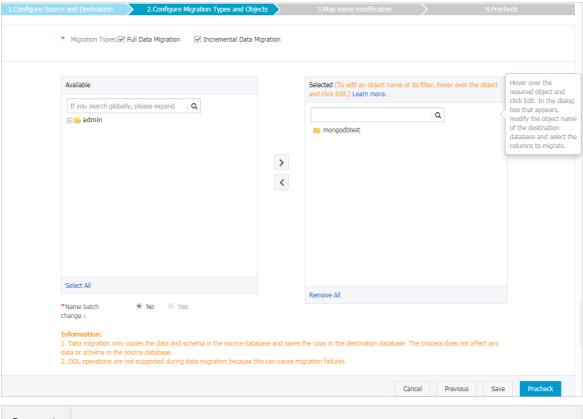
Instance Region	If Instance Type is set to User-Created Database with Public IP Address, you do not need to specify the Instance Region. Onte If you have configured a whitelist for the user-created MongoDB database, you must add the CIDR blocks of DTS servers to the whitelist. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.	
Database Type	Select MongoDB.	
Hostname or IP Address	Enter the endpoint of a shard for the source database. In this example, enter the public IP address of the shard. Onte DTS migrates each shard of the source database in turn. In this example, enter the endpoint of the first shard. Then enter the endpoint of the second migration task. Repeat this operation until all shards are migrated.	
Port Number	Enter the service port of the shard. Note The service port of each shard for user-created MongoDB databases must be open to the public network.	
Dat abase Name	Enter the name of the authentication database to which the database account belongs.	
Dat abase Account	Enter the username of the database account that you use to manage the source database.	
Database Password	Pote After you specify the source database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the Failed message is displayed, click Check in the Failed message to modify the information as prompted.	
Connection Method	Select Non-encrypted. Note The SSL-encrypted option is available only when you migrate MongoDB Atlas.	
	Region Database Type Hostname or IP Address Port Number Database Name Database Account Database Account	

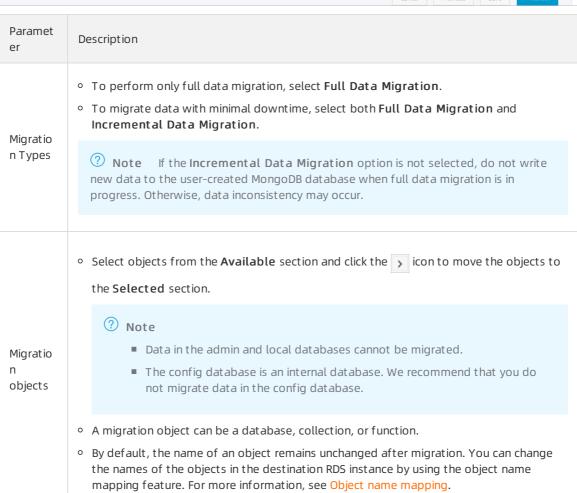
Section	Parameter	Description		
	Instance Type	Select MongoDB Instance.		
	Instance Region	Select the region in which the ApsaraDB for MongoDB instance resides.		
	MongoDB Instance ID	Select the ID of the ApsaraDB for MongoDB instance.		
	Database Name	Enter the name of the authentication database to which the database account belongs.		
Destinatio n		Note If you want to use the root account, specify admin for the Database Name parameter.		
Database	Dat abase Account	Enter the username of the database account that you use to manage the source database.		
		Enter the password for the destination database account.		
	Database Password	Note After you specify the destination database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed, and you must click Check next to the Failed message to modify the information as prompted.		

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

Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination RDS instance. This ensures that DTS servers can connect to the destination ApsaraDB for MongoDB instance. After the migration is completed, you can remove these CIDR blocks from the whitelist. For more information, see Configure a whitelist for a sharded cluster instance.

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

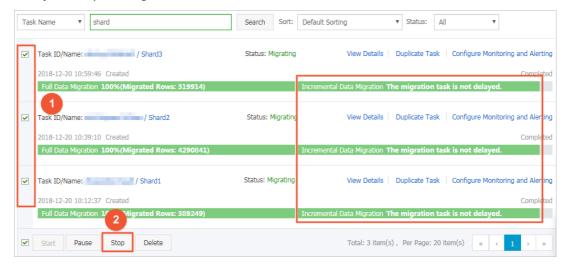




8. In the lower-right corner of the page, click **Precheck**.



- 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.
- o If the task fails to pass the precheck, click the icon next to each failed item to view details. Troubleshoot the issues based on the causes 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 and select Data Transmission Service (Pay-As-You-Go) Service Terms.
- 11. Click Buy and Start to start the migration task.
- 12. Repeat step 1 to step 11 to create data migration tasks for the remaining shards.
- 13. Stop the data migration task.
 - Full data migration
 Do not manually stop a task during full data migration. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task automatically stops.
 - Incremental data migration
 An incremental data migration task does not automatically stop. You must manually stop the migration task.
 - **? Note** Select an appropriate time to manually stop the migration task. For example, you can stop the migration task during off-peak hours or before you switch your workloads to the destination instance.
 - 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. After the status of **Incremental Data Migration** changes to **The migration task** is **not delayed**, stop the migration task.



14. Switch your workloads to the ApsaraDB for MongoDB instance.

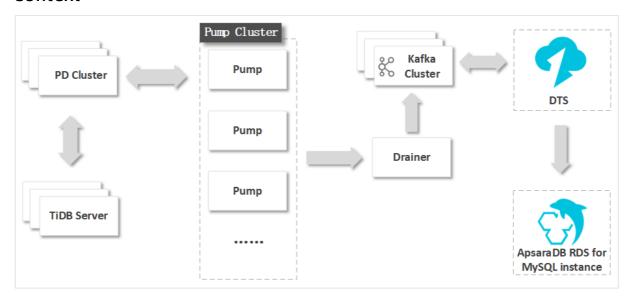
7.7. Source: TiDB

7.7.1. Migrate incremental data from a user-created TiDB database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate incremental data from a user-created 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

Context



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 Alibaba Cloud over the Internet. For more
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

Migration types

Migration type	Description	
	DTS migrates the schemas of the required objects to the destination database. DTS supports schema migration for views, tables, and databases.	
Schema migration	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	
	DTS migrates historical data of the required objects to the destination database.	
Full data migration	Note 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.	
Incremental data	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: • Data manipulation language (DML) operations: INSERT, UPDATE, and DELETE	
migration	Data definition language (DDL) operations: CREATE TABLE, DROP TABLE, ALTER TABLE, RENAME TABLE, TRUNCATE TABLE, CREATE VIEW, DROP VIEW, and ALTER VIEW	
	Incremental data migration allows you to ensure service continuity when you migrate data from a user-created TiDB database to Alibaba Cloud.	

Before you begin

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 incremental data migration.

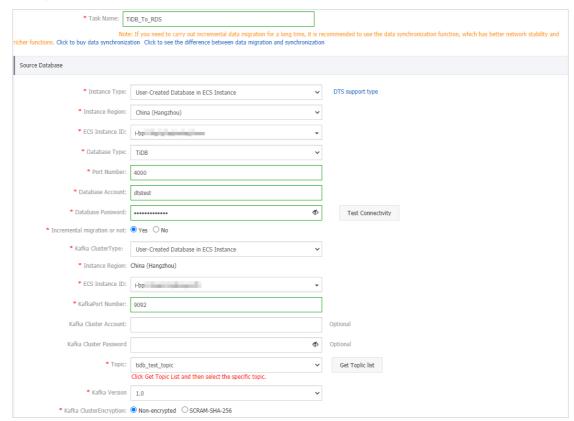
- 1. Deploy Pump and Drainer. For more information, visit TiDB Binlog Cluster Deployment.
- 2. Modify the configuration file of Drainer and specify a Kafka cluster to receive data from Drainer. For more information, visit Binlog Slave Client User Guide.
- 3. Deploy a Kafka cluster by using one of the following methods:
 - Deploy a user-created Kafka cluster. For more information, visit Apache Kafka official website.

Warning We recommend that you set the message.max.bytes and replica.fetch. max.bytes parameters for the Kafka broker to greater values. We also recommend that you set the fetch.message.max.bytes parameter for the Kafka consumer to a greater value. These settings ensure that the Kafka cluster can receive the binary log files that are generated in TiDB. For more information, visit Kafka 2.5 Documentation.

- Purchase and deploy a Message Queue for Apache Kafka instance. For more information, see Quick start in the Message Queue for Apache Kafka documentation.
 - **? Note** The Message Queue for Apache Kafka instance must be deployed in the same 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 user-created Kafka cluster or the Message Queue for Apache Kafka instance.
- 5. Add the CIDR blocks of DTS servers to the 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 for the data migration task.
 - i. Configure the task name and source database.

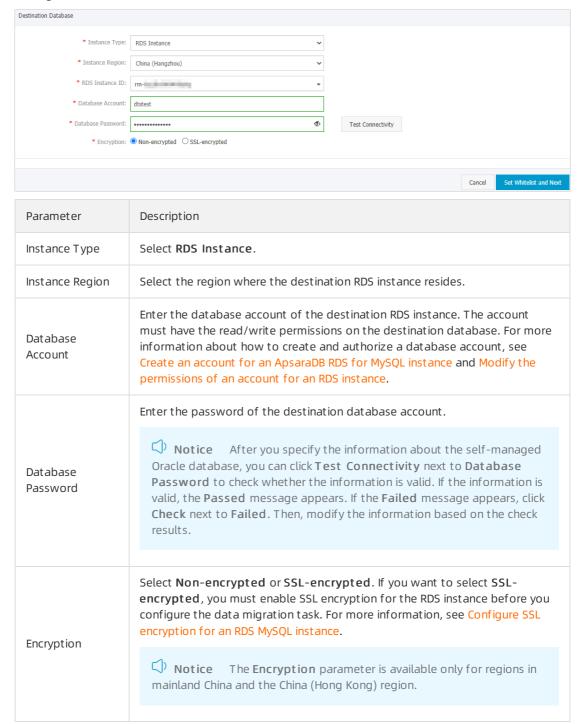


Parameter	Description
Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
Instance Type	Select an instance type based on where the source database is deployed. User-Created Database in ECS Instance is used in this example.
	Note 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	Select the region of the ECS instance on which the source TiDB database is deployed.
Database Type	Select TiDB .

Parameter	Description
Port Number	Enter the service port number of the source TiDB database. The default port number is 4000 .
Database Account	Enter the account of the source TiDB database. The account must have the SELECT permission on the objects to be migrated and the SHOW VIEW permission.
	Enter the password of the source database account.
Database Password	Notice 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.
Incremental migration or not	Specify whether to perform incremental data migration. Yes is selected in this example. For 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	Select an instance type based on where the user-created Kafka cluster is deployed. User-Created Database with Public IP Address is selected in this example.
	Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview.
Instance Region	The value of this parameter is the same as the region of the source database and cannot be changed.
ECS Instance ID	Select the ID of the ECS instance on which the user-created Kafka cluster is deployed.
Kafka Port Number	Enter the service port number of the user-created Kafka cluster. The default port number is 9092.
Kafka Cluster Account	Enter the username that is used to log on to the user-created Kafka cluster. If authentication is disabled for the Kafka cluster, you do not need to enter the username.
Kafka Cluster Password	Enter the password of the cluster account. If authentication is disabled for the Kafka cluster, you do not need to enter the password.
Topic	Click Get Topic List , and select a topic name from the drop-down list.
Kafka Version	Select the version of the user-created Kafka cluster.

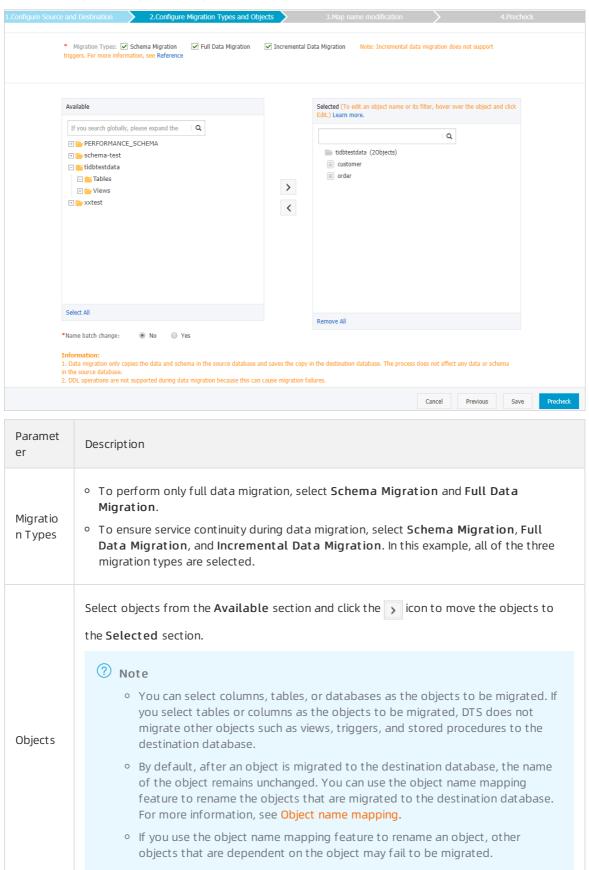
Parameter	Description
Kafka Cluster Encryption	Select Non-encrypted or SCRAM-SHA-256 based on your business and security requirements.

ii. Configure the destination database.



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

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



Paramet er Description

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



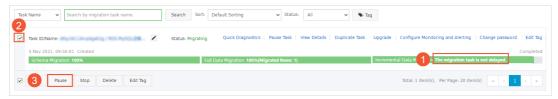
- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- \circ If the task fails to pass the precheck, you can click the \bigcirc icon next to each failed item

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.
- 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.



7.7.2. Migrate full data from a user-created TiDB database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a user-created TiDB database to an ApsaraDB RDS for MySQL instance by using Data Transmission Service (DTS).

Prerequisites

An 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 reside in the China (Hangzhou), China (Shanghai), China (Qingdao), China (Beijing), China (Shenzhen), China (Zhangjiakou-Beijing Winter Olympics), China (Hong Kong), Singapore, US (Silicon Valley), US (Virginia). The feature is available only in these regions.
- The available storage space of the destination ApsaraDB RDS for MySQL instance is larger than the total size of the data in the user-created 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	Instance configuration	Internet traffic
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 the required objects to the destination database. DTS supports schema migration for the following types of objects: database, table, and view.		
Schema migration	warning Heterogeneous databases support different data types. DTS maps a data type in the source database to a data type supported by the destination database during schema migration. For more information, see Data type mappings between heterogeneous databases.		
	DTS migrates historical data of the required objects to the destination database.		
Full data migration	Notice 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.		

Permissions required for database accounts

Database	Schema migration	Full data migration
TiDB database	The SELECT permission	The SHOW VIEW permission and the SELECT permission on the required objects.
ApsaraDB RDS for MySQL instance	The read/write permissions	The read/write permissions

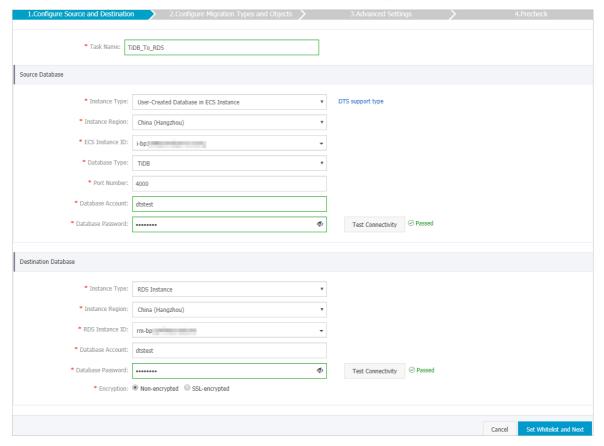
For more information about how to create and authorize a database account in the ApsaraDB RDS for MySQL instance, see Create an account for an ApsaraDB RDS for MySQL instance and Modify the permissions of an account for an RDS instance.

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 for the data migration task.

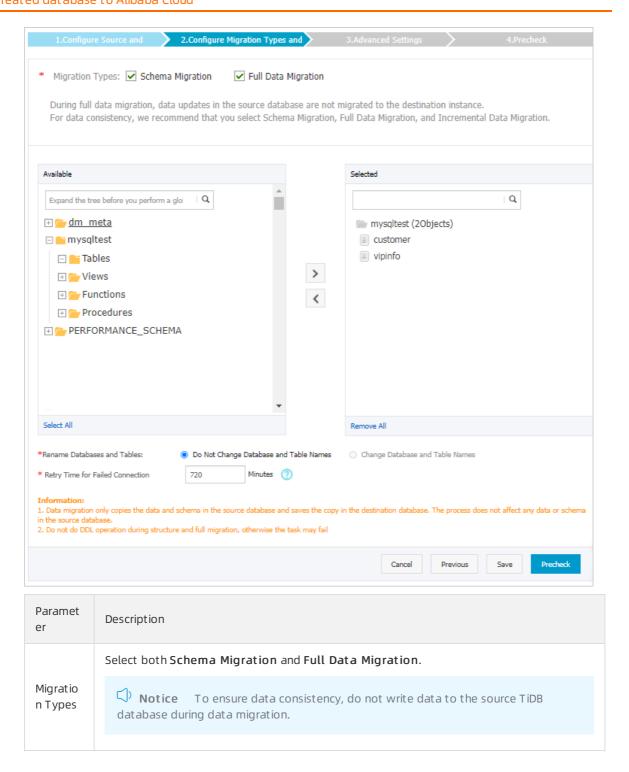


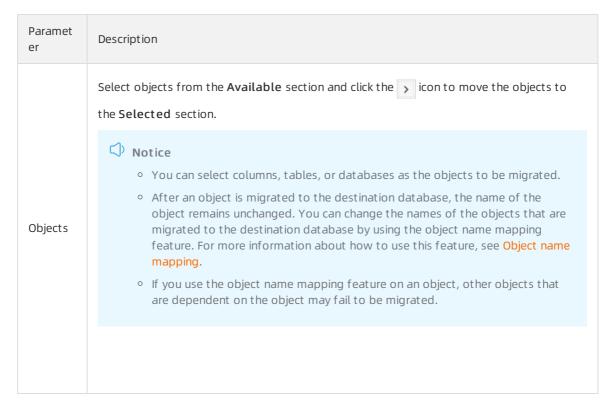
Parameter	Description		
Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.		
Instance Type	Select an instance type based on where the source database is deployed. The procedure in this topic uses a user-created database hosted on ECS as an example. Note 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	Select the region of the ECS instance on which the source TiDB database is deployed.		
Database Type	Select TiDB.		
Port Number	Enter the service port number of the source TiDB database. The default port number is 4000 .		
Dat abase Account	Enter the account of the source TiDB database.		
Dat abase Password	Enter the password of the source database account. Notice After you specify the information about the selfmanaged 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 RDS Instance.		
Instance Region	Select the region where the destination RDS instance resides.		
RDS Instance ID	Select the ID of the destination RDS instance.		
	Instance Type Instance Region Database Type Port Number Database Account Database Password		

Section	Parameter	Description		
		Enter the password of the destination database account.		
Destinatio n Database	Dat abase Password	Notice 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.		
	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 for an RDS MySQL instance.		
		Notice The Encryption parameter is available only for regions in mainland China and the Hong Kong (China) region.		

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

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





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



- 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.
- \circ If the task fails to pass the precheck, you can click the $_{\boxed{\emph{0}}}$ 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.
- 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 migration task.

• Warning We recommend that you do not manually stop a data migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the migration task automatically stops.



7.8. Source: Db2

7.8.1. Migrate data from a user-created Db2 database to an ApsaraDB RDS for MySQL instance

This topic describes how to migrate data from a user-created 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 user-created Db2 database, you can select all of the supported migration types to ensure service continuity.

Prerequisites

- The Db2 database is deployed on Linux.
- The version of the Db2 database is 9.7 or 10.5.
- The available storage space of the destination ApsaraDB RDS for MySQL instance is larger than the total size of the data in the user-created Db2 database.

Precautions

- DTS does not synchronize data definition language (DDL) operations to the destination instance.
- 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 databases and accounts.

- 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.

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
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 dat a migration
 - DTS migrates historical data of the required objects from the source Db2 database to the destination ApsaraDB RDS for MySQL instance.
- Increment al dat a migration

After full data migration is complete, DTS synchronizes incremental data from the source Db2 database to the destination ApsaraDB RDS for MySQL instance. Incremental data migration allows you to ensure service continuity when you migrate data from a user-created 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/write permissions	The read/write permissions	The read/write permissions

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

- Creating group and user IDs for a Db2 database installation (Linux and UNIX) and Authorities overview for a user-created Db2 database
- Create an account for an ApsaraDB RDS for MySQL instance and Modify the permissions of an account for an RDS instance for an ApsaraDB RDS for MySQL instance

Data migration process

To avoid data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source 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 increment al data migration.

Before you begin

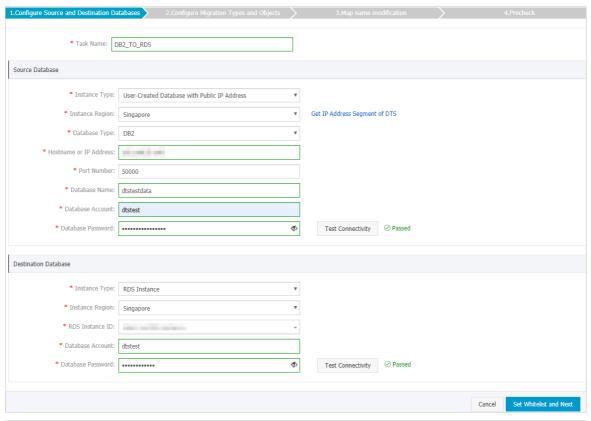
Before you configure an incremental data migration task, enable the archive log feature for the source Db2 database. For more information, see Primary log archive method and Secondary log archive method.

? 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 for the data migration task.



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 where the source database is deployed. In this example, select User-Created Database with Public IP Address .	
	Instance Type	Note If you select other instance types, you must deploy the network environment for the self-managed database. For more information, see Preparation overview.	
		If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.	
	Instance Region	Note If a whitelist is configured for the user-created 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.	

Section	Parameter	Description	
Source Dat abase			
	Database Type	Select DB2.	
	Hostname or IP Address	Enter the endpoint that is used to connect to the user-created Db2 database. In this example, enter the public IP address.	
		Enter the service port number of the user-created Db2 database. The default port number is 50000 .	
	Port	Note The service port of the user-created Db2 database must be accessible over the Internet.	
	Dat abase Name	Enter the name of the user-created Db2 database.	
	Dat abase Account	Enter the account of the user-created Db2 database.	
		Enter the password of the source 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 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.	

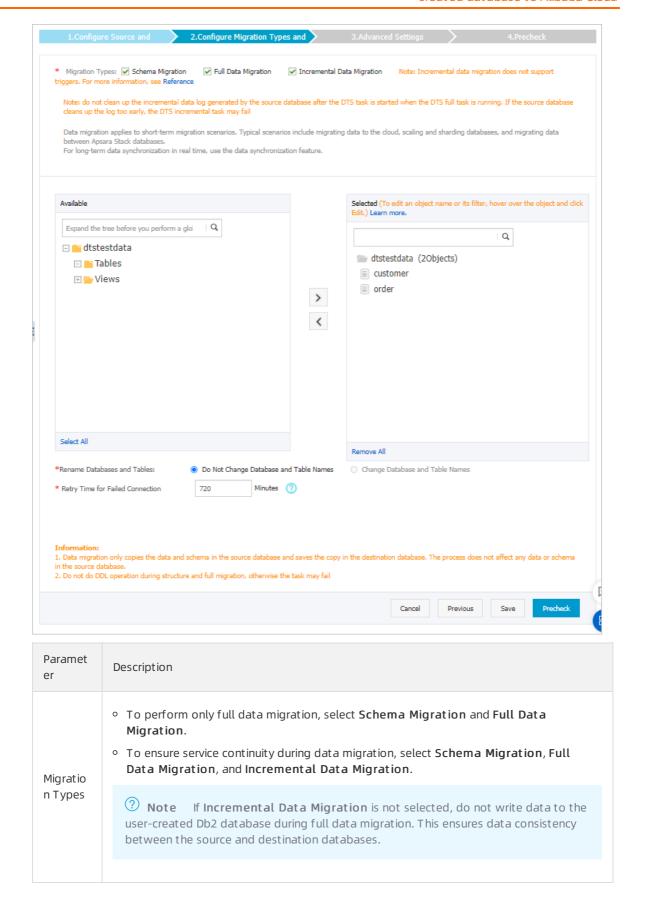
Section	Parameter	Description		
Destinatio		Enter the password of the destination database account.		
n Database	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the 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 . 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 for an RDS MySQL instance.		
		Note The Encryption parameter is available only for regions in mainland China and the China (Hong Kong) region.		

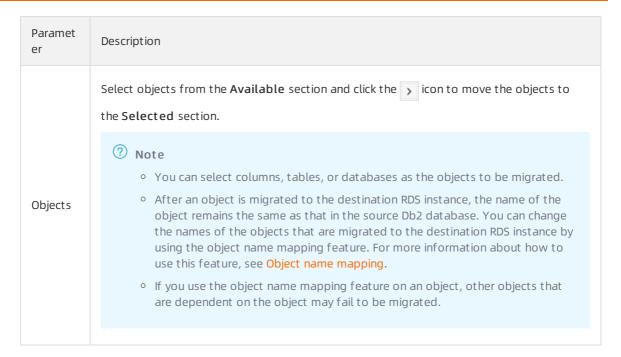
6. In the lower-right corner of the page, click **Set Whitelist and Next**.

? 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 objects to be migrated.

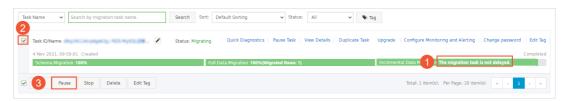
195





8.

- 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 self-managed 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.

8.Migrate from a third-party cloud to Alibaba Cloud

8.1. 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. 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**. This 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

199

- 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
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

- o During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.
- o DTS does not migrate user information. Before a user can call views, stored procedures, and functions of the destination database, you must grant the read/write permissions to the user.
- Full dat a migration

DTS migrates historical data of the required objects from the Amazon RDS for MySQL instance to the ApsaraDB RDS for MySQL instance.



? Note

- o 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.
- o During schema migration and full data migration, do not perform 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

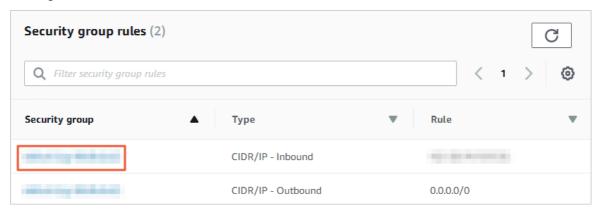
Database	Schema migration	Full data migration	Incremental data migration
Amazon RDS for MySQL	The SELECT permission	The SELECT permission	The REPLICATION CLIENT, REPLICATION SLAVE, SHOW VIEW, and SELECT permissions
ApsaraDB RDS for MySQL	The read/write permissions	The read/write permissions	The read/write permissions

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

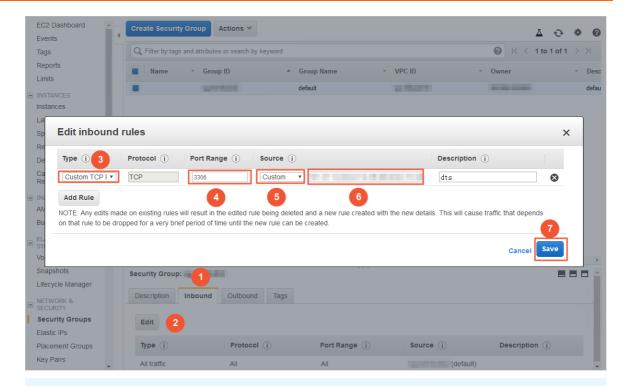
- Create an account for a user-created MySQL database and configure binary logging for an Amazon RDS for MySQL instance
- Create an account for an ApsaraDB RDS for MySQL instance and Modify the permissions of an account for an RDS instance for an ApsaraDB RDS for MySQL instance

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.



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.



? Note

- You need to add only the CIDR blocks of DTS servers that are located in the same region
 as the destination database. For example, the source database is located in Singapore
 and the destination database is located in Hangzhou. You need to add only the CIDR
 blocks of DTS servers that are located 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.

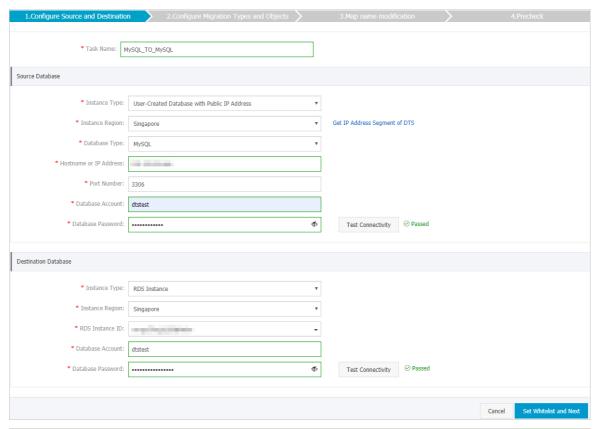
call mysgl.rds set configuration('binlog retention hours', 24);

? Note

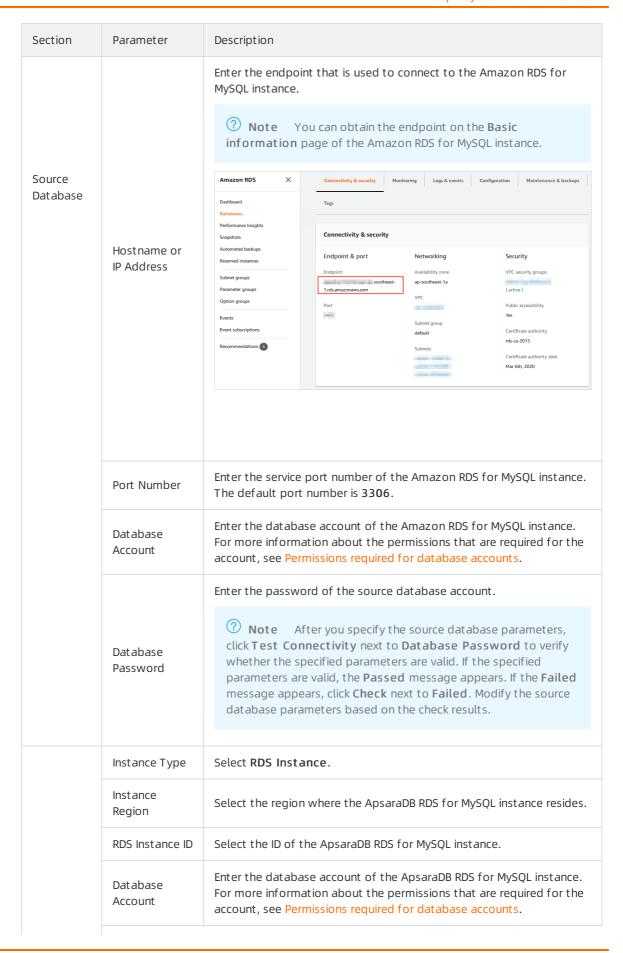
- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of the Amazon RDS for MySQL instance must be enabled and the value of the binlog_format parameter must be set to row. If the version of MySQL is
 5.6 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 for the data migration task.



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	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.
	Database Type	Select MySQL.

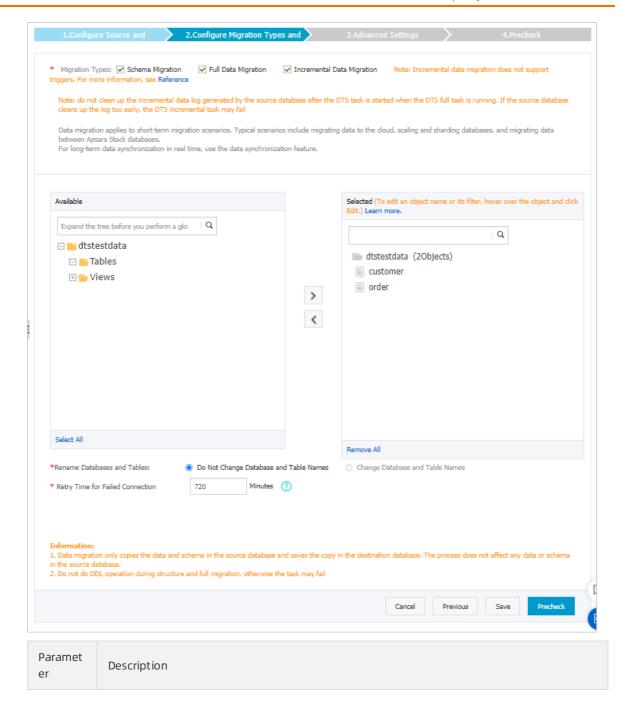


Section	Parameter	Description	
		Enter the password of the destination database account.	
Destinatio n Database	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the specified parameters are valid. If the 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. 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 for an RDS MySQL instance.	
		Note 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**.

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

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



Paramet er	Description
Migratio n Types	 To perform only full data migration, select Schema Migration and Full Data Migration. 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, do not write data to the source database during full data migration. This ensures data consistency between the source and destination databases. 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.
	Select objects from the Available section and click the icon to move the objects to the Selected section.
Objects	 Note You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains unchanged. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.

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



- 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.

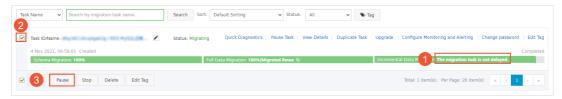
- You can troubleshoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 9. After the task passes the precheck, click **Next**.
- 10. In the Confirm Settings dialog box, specify the Channel 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.

8.2. Migrate data from an Amazon RDS for Oracle database to an ApsaraDB RDS for MySQL database

This topic describes how to migrate data from an Amazon RDS for Oracle database to an ApsaraDB RDS for MySQL database by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. You can select all of the three migration types to migrate data without service interruptions.

Prerequisites

- **Public Availability** is set to **Yes** for the Amazon RDS for Oracle database. This ensures that DTS can access the Amazon RDS for Oracle database through the public network.
- The version of the Amazon RDS for Oracle database is 9i, 10g, or 11g.
- The version of the ApsaraDB RDS for MySQL database is 5.6 or 5.7.
- The available storage space of the ApsaraDB RDS for MySQL database is larger than the total size of the data in the Amazon RDS for Oracle database.

Note The binary log files generated during the migration occupy some space. They are automatically cleared when the 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 Alibaba Cloud over the Internet. For more
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

Migration types

- Schema migration
 DTS supports schema migration of tables, indexes, constraints, and sequences. However, DTS does
 not support schema migration of views, synonyms, triggers, stored procedures, stored functions,
 packages, and custom object types.
- Full dat a migration

DTS migrates full data of objects to be migrated from the Amazon RDS for Oracle database to the database in the ApsaraDB RDS for MySQL instance.

Incremental data migration
 During full data migration, DTS polls and captures the redo log files from the Amazon RDS for Oracle database. After the full data migration is complete, DTS synchronizes incremental data from the Amazon RDS for Oracle database to the destination ApsaraDB RDS for MySQL database. Incremental data migration helps you achieve service continuity when migrating data from an Amazon RDS for Oracle database to the destination database.

SQL operations that can be synchronized during incremental data migration

- INSERT, DELETE, and UPDATE
- CREATE TABLE

Note The operations on tables or partitions that contain computed columns cannot be synchronized.

- ALTER TABLE operations, including 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 database	Owner permissions on schemas to be migrated	Owner permissions on schemas to be migrated	Master user permissions
ApsaraDB RDS for MySQL.	Read and write permissions on the destination database	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:

- CREATE USER and GRANT for an Amazon RDS for Oracle database
- Create accounts and databases for an RDS for MySQL instance

Data type mappings

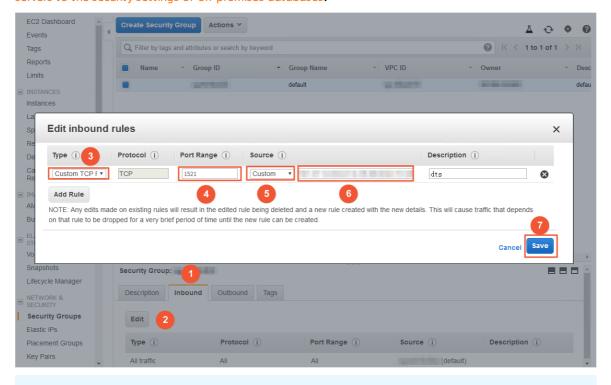
For more information, see Data type mappings between heterogeneous databases.

Preparation

- 1. Log on to the Amazon RDS Management Console.
- 2. Go to the **Basic Information** page of the source 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.



4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit, and add CIDR blocks of DTS servers to the Edit inbound rules dialog box that appears. The CIDR blocks of DTS servers vary with the region where the destination database resides. For more information about the CIDR blocks of DTS servers, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.



- ? 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. Change log configurations for the Amazon RDS for Oracle database. Skip this step if you do not need to perform incremental data migration.
 - i. Use the master user account and the SQL*Plus tool provided by Oracle to connect to the

Amazon RDS for Oracle database.

- ii. Run the archive log list; command to check whether archiving is enabled for the Amazon RDS for Oracle database.
 - **Note** If archiving is disabled for the instance, enable the archiving mode. For more information, see Managing Archived Redo Logs.
- iii. Enable force logging.

```
exec rdsadmin.rdsadmin_util.force_logging(p_enable => true);
```

iv. Enable supplemental logging for primary keys.

```
begin rdsadmin.rdsadmin_util.alter_supplemental_logging(p_action => 'ADD',p_type =>
'PRIMARY KEY');end;/
```

v. Enable supplemental logging for unique keys.

```
begin rdsadmin.rdsadmin_util.alter_supplemental_logging(p_action => 'ADD',p_type =>
'UNIQUE');end;/
```

vi. Set a retention period for archived logs.

```
begin rdsadmin.rdsadmin_util.set_configuration(name => 'archivelog retention hours'
, value => '24');end;/
```

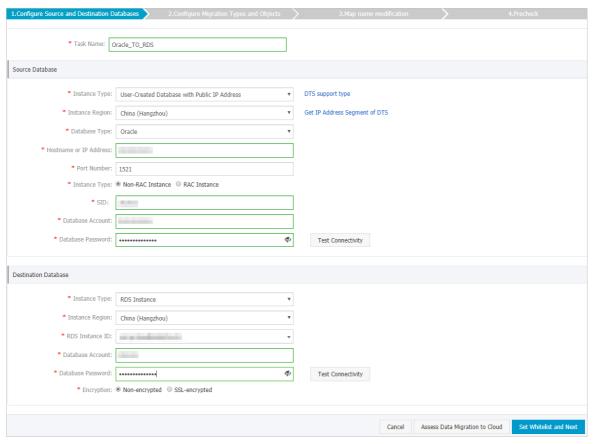
Note We recommend that you set the retention period of archived logs to at least 24 hours.

vii. Commit the changes.

```
commit;
```

Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Dat a 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 generates a random task name. However, we recommend that you specify an informative name to ease management.
	Instance Type	Select User-Created Database with Public IP Address.
	Instance Region	When the instance type is set to User-Created Database with Public IP Address , you do not need to set Instance Region .
	Database Type	Select Oracle.

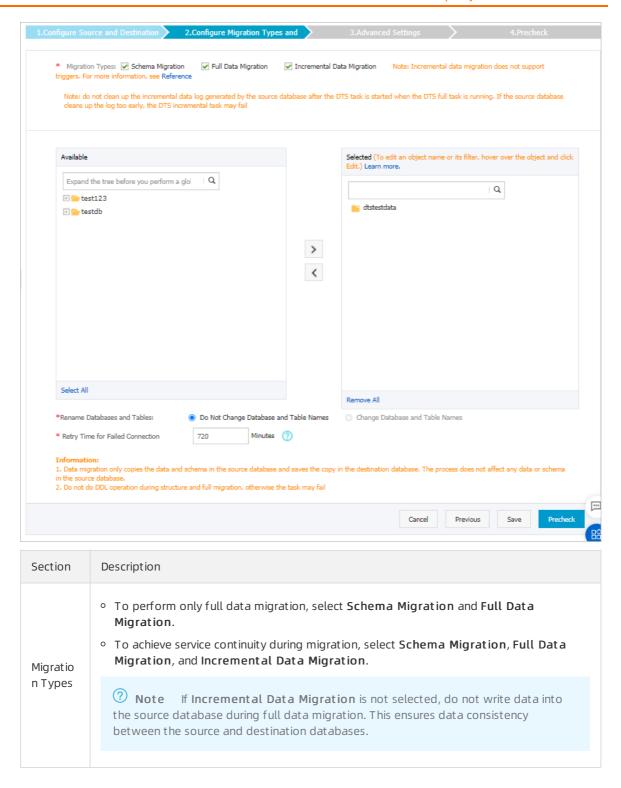
Section	Parameter	Description		
		Enter the endpoint for accessing the Amazon RDS for Oracle database. Note You can obtain the endpoint on the Basic information page of the source Amazon RDS for Oracle database.		
Database	Hostname or IP Address	Amazon RDS X Connectivity & security Monitoring Logs & events Configuration Maintenance & backups Dashboard Tags Databases Performance insights Snapshots Connectivity & security		
		Automated backups Reserved instances Subnet groups Parameter groups Option groups Events Event subscriptions Recommendations (3) Endpoint & port Endpoint Availability zone apsoutheast-1a (active) VPC Public accessibility Yes Subnet group default Certificate authority rds-ca-2015 Subnets Certificate authority date Mar 6th, 2020		
	Port Number	Enter the port number configured for the Amazon RDS for Oracle database. The default port number is 1521 .		
Instance Type		 Non-RAC Instance: If you select this option, you must specify the SID. RAC Instance: If you select this option, you must specify the Service Name. In this example, select Non-RAC Instance and enter SID. 		
Dat abase Account		Enter the account for the Amazon RDS for Oracle database. For more information about permissions required for the account, see Permissions required for database accounts.		
	Dat abase Password	Pote 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 Instance.		
	Instance Region	Select the region where the destination ApsaraDB RDS for MySQL instance resides.		

Section	Parameter	Description	
	RDS Instance ID	Select the RDS instance ID.	
Destinatio n Database	Dat abase Account	Enter the database account of the ApsaraDB RDS for MySQL instance. For more information about permissions 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.

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.
- o 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. Configure migration types and objects.



Section	Description
Objects to be migrate d	Select the objects to be migrated in the Available section and click to move them to the Selected section. Note Objects to be migrated can be databases, tables, or columns. By default, the selected objects are not renamed after the migration. If you want to rename the objects that are migrated to the destination instance, you can use the object name mapping feature provided by DTS. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature for an object, objects that depend
	on the object may fail to be migrated.

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



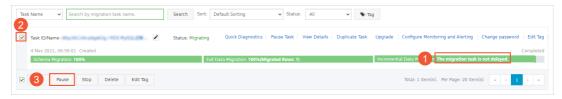
- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- \circ If the task fails to pass the precheck, you can click the \bigcirc icon next to each failed item

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.
- 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 database.

8.3. Migrate incremental data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL database

This topic describes how to migrate incremental data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL database 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 database is 10.4-R1 to 11.4-R1.
- The **Public accessibility** option of Amazon RDS for PostgreSQL is set to **Yes**. This ensures that DTS can access Amazon RDS for PostgreSQL over the Internet.
- The value of the rds.logical_replication parameter is set to 1. This ensures that DTS can read incremental data from the Amazon RDS for PostgreSQL database.
- An ApsaraDB RDS for PostgreSQL instance is created. For more information, see Create an RDS for PostgreSQL instance.



- The version of the ApsaraDB RDS for PostgreSQL database 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 database is larger than the total size of the data in the Amazon RDS for PostgreSQL database.

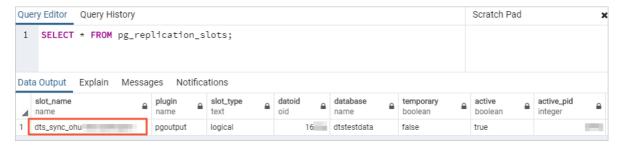
Notes

• 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, duplicate data may exist in the destination database and 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.
- Only data manipulation language (DML) operations, such as INSERT, UPDATE, DELETE, and UPDATE operations, can be migrated.
- During data migration, DTS creates a replication slot for the Amazon RDS for PostgreSQL database. The replication slot is prefixed with dts_sync_ . DTS automatically clears historical replication slots every 90 minutes to prevent them from occupying disk space.

If a migration task is released or fails, DTS automatically clears the replication slot. After a switchover between primary and secondary Amazon RDS for PostgreSQL databases, you must log on to the secondary database to manually clear the replication slot.



• DTS automatically resumes a failed data migration task. Before switching your workloads to the destination instance, you must stop or release the data migration task. Otherwise, the data from the source database 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
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 database. DTS supports schema migration for the following types of objects: table, trigger, view, sequence, function, user-defined type, rule, domain, operation, and aggregate.

? Note Functions that are written in the C programming language cannot be migrated.

- Full data migration
 DTS migrates historical data of the required objects from the Amazon RDS for PostgreSQL database to the ApsaraDB RDS for PostgreSQL database.
- Increment al data migration
 After full data migration is complete, DTS synchronizes incremental data from the Amazon RDS for
 PostgreSQL database to the ApsaraDB RDS for PostgreSQL database. 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 database	The usage permission for pg_catalog	The SELECT permission for the objects to be migrated	The rds_superuser permission
ApsaraDB RDS for PostgreSQL database	The create and usage permissions for the objects to be migrated	The owner permission for schemas	The owner permission for schemas

Data migration process

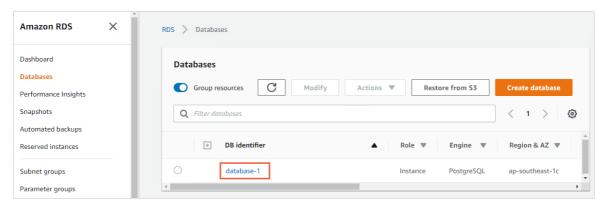
To avoid data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

- 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 performing incremental data migration, do not perform DDL operations on the objects to be migrated in the Amazon RDS for PostgreSQL database. Otherwise, the objects may fail to be migrated.

Preparations

- 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.



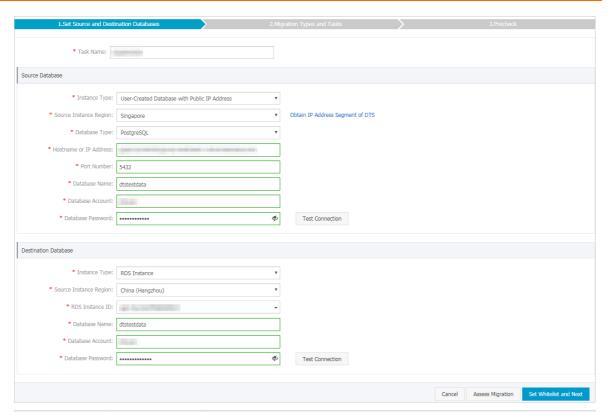
4. In the **Security group rules** section, click the name of the security group corresponding to the existing inbound rule.



5.

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 information about the source and destination databases for the data migration task.



Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you use 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 User-Created Database with Public IP Address , you do not need to specify the instance region.
	Dat abase Type	Select PostgreSQL.

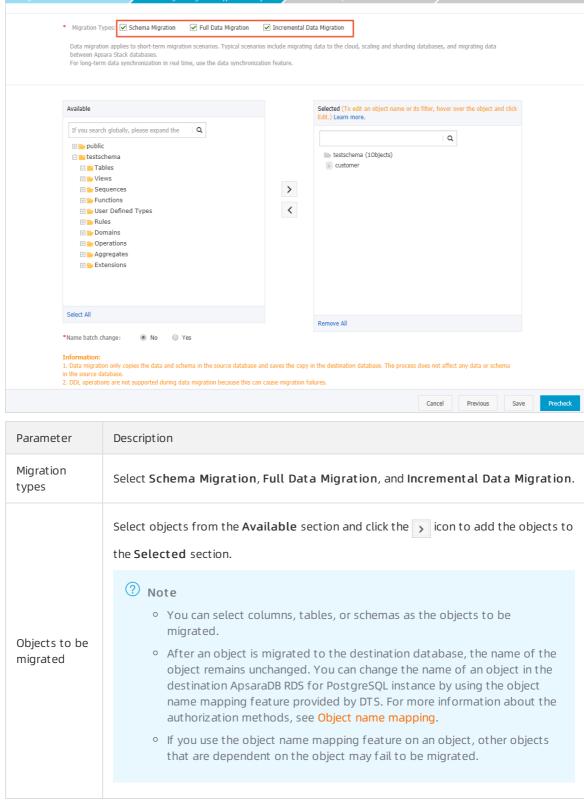
Hostname or IP Address Port Number Enter the service port number of the Amazon RDS for PostgreSQL database. The default port number is 5432. Database Name Database Account Database Password Databas	Section	Parameter	Description		
Port Number Enter the service port number of the Amazon RDS for PostgreSQL database. The default port number is 5432. Database Name Enter the name of the source Amazon RDS for PostgreSQL database. For more information about permissions required for the account, see Permissions required for database accounts. Enter the password for the database accounts. Enter the password for the database accounts. 2 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 Instance. Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.			Enter the endpoint that is used to connect to the Amazon RDS for PostgreSQL database. Note You can find the endpoint on the Basic information page of the source Amazon RDS for PostgreSQL instance. Amazon RDS X Connectivity & security Tags Connectivity & security Tags Connectivity & security Networking Reserved instances Subnet groups Parameter groups Option groups Everts Ev		
Database Account Database Permissions required for the account, see Permissions required for database accounts. Enter the password for the database account. Onte 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 Instance. Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.		Database	Enter the service port number of the Amazon RDS for PostgreSQL		
Database Password Password Password Database Password Password Database Password Database Password Database Password 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 Instance. Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.			Enter the account for the Amazon RDS for PostgreSQL database. For more information about permissions required for the account, see Permissions required for database accounts.		
Instance Select the region where the destination ApsaraDB RDS for Region PostgreSQL instance resides.		20100000	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		
Region PostgreSQL instance resides.		Instance Type	Select RDS Instance.		
DDC Instance Colors the ID of the destination AnswerD DDC for Doctor-COL					
ID select the iD of the destination Apsarabb RDS for PostgreSQL instance.		RDS Instance	Select the ID of the destination ApsaraDB RDS for PostgreSQL instance.		

Section	Parameter	Description	
	Dat abase Name	Enter the name of the destination ApsaraDB RDS for PostgreSQL database. The name can be different from the name of the source Amazon RDS for PostgreSQL database.	
Destination Database		Note Before you configure the data migration task, you must create a database in the ApsaraDB RDS for PostgreSQL instance. For more information, see Create a database.	
	Dat abase Account	Enter the database account for the destination ApsaraDB RDS for PostgreSQL instance. For more information about permissions required for the account, see Permissions required for database accounts.	
	Dat abase Password	Enter the password for 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.	

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

Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination ApsaraDB RDS for PostgreSQL instance. This ensures that DTS servers can connect to the destination ApsaraDB RDS for PostgreSQL instance.

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



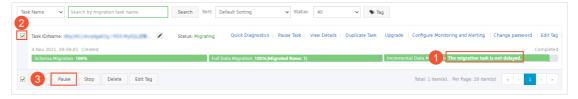
8. In the lower-right corner of the page, click **Precheck**.



- 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.

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 migration task.
 - **Note** The migration task does not automatically end. You must manually stop the migration task. Select an appropriate time to manually stop the migration task. For example, you can stop the migration 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 appears 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. After the status of incremental data migration changes to The migration task is not delayed, manually stop the migration task.



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

8.4. Migrate full data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL database

This topic describes how to migrate full data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL database by using Data Transmission Service (DTS).

Prerequisites

- The version of the Amazon RDS for Post greSQL database is 9.4, 9.5, 9.6, or 10.0.
- The **Public accessibility** option of Amazon RDS for PostgreSQL is set to **Yes**. This ensures that DTS can access Amazon RDS for PostgreSQL over the Internet.
- An ApsaraDB RDS for PostgreSQL instance is created. For more information, see Create an RDS for PostgreSQL instance.
 - ? Note The version of the ApsaraDB RDS for PostgreSQL database is 9.4 or 10.0.
- The available storage space of the ApsaraDB RDS for PostgreSQL instance is larger than the total space of the data in the Amazon RDS for PostgreSQL instance.

Notes

- 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.
- 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.
- DTS automatically resumes a failed data migration task. Before switching your workloads to the destination instance, you must stop or release the data migration task. Otherwise, the data from the source database will overwrite the data in the destination instance after the task is resumed.

Billing

Migration type	Instance configurations	Internet traffic
Schema migration or full data migration	Free of charge	Migrating data from Alibaba Cloud over the Internet incurs fees. For more information, see Pricing.

Migration types

- Schema migration
 - DTS migrates the schemas of the required objects to the ApsaraDB RDS for PostgreSQL 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 Amazon RDS for PostgreSQL database to the ApsaraDB RDS for PostgreSQL database.

Permissions required for database accounts

Database	Schema migration	Full data migration
Amazon RDS for PostgreSQL database	The usage permission for pg_catalog	The permission to perform SELECT operations on the objects to be migrated

Database	Schema migration	Full data migration
ApsaraDB RDS for PostgreSQL database	The create and usage permissions on the objects to be migrated	The owner permission for schemas

Full data migration

To avoid data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

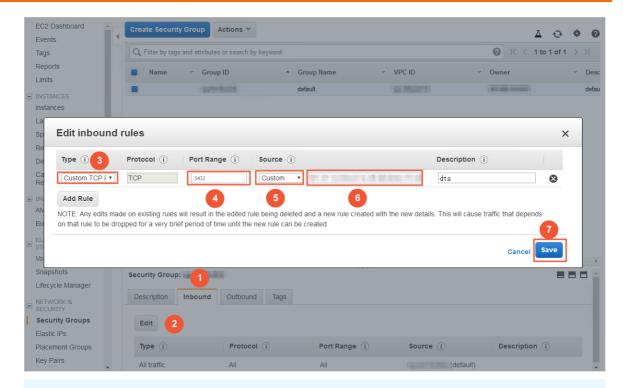
- 1. Migrate the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates.
- 2. Perform full data migration.
- 3. Migrate the schemas of triggers and foreign keys.

Preparation 1: Modify Amazon inbound rules

- 1. Log on to the Amazon RDS Management Console.
- 2. Go to the **Basic Information** page of the source 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.



4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit, and add CIDR blocks of DTS servers to the Edit inbound rules dialog box that appears. The CIDR blocks of DTS servers vary with the region where the destination database resides. For more information about the CIDR blocks of DTS servers, see Add the CIDR blocks of DTS servers to the security settings of on-premises databases.



? Note

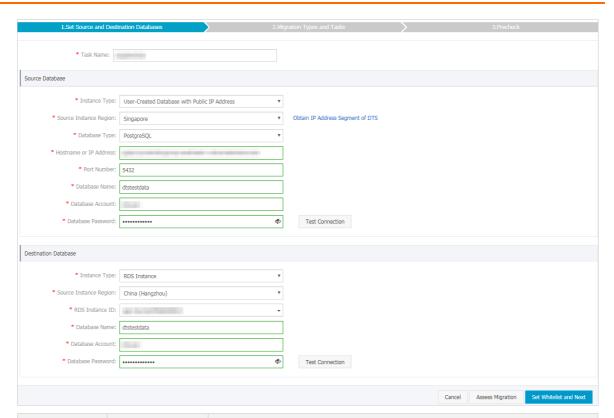
- You need to add only the CIDR blocks of DTS servers that reside in the same region as the destination database. For example, the source database resides in the Singapore (Singapore) region and the destination database resides in the China (Hangzhou) region.
 You need to add only the CIDR blocks of 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 destination database must be the same as that of the source database. For more information, see Create a database 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 $\mbox{\it Create Migration Task}.$
- 5. Configure the information about the source and destination databases for the data migration task.



Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you use 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 User-Created Database with Public IP Address , you do not need to specify the instance region.
	Database Type	Select PostgreSQL.

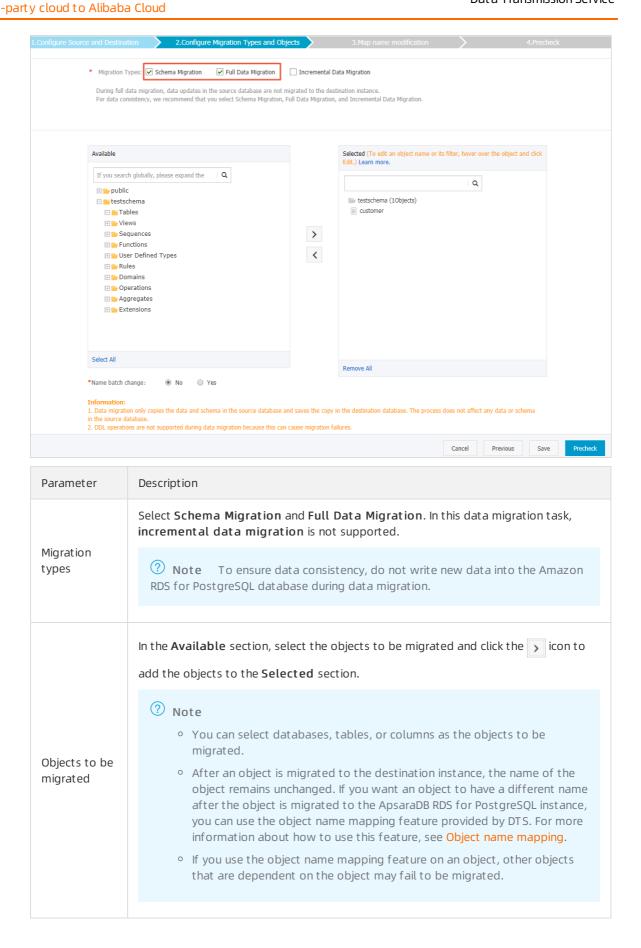
Section	Parameter	Description		
Source Dat abase	Hostname or IP Address	Enter the endpoint that is used to connect to the Amazon RDS for PostgreSQL database. To Note You can obtain the endpoint on the Basic information page of the source Amazon RDS for PostgreSQL instance. Amazon RDS X Connectivity & security Monitoring Logs & events Configuration Maintenance & backups Tags Connectivity & security Networking Security Endpoint & Port Public accessibility Ferents Subnets groups Connectivity & southeast- 1 Little Amazon RDS Configuration Maintenance & backups Networking Security Port Public accessibility Yes Subnets Certificate authority date Mar 6th, 2020		
	Port Number	Enter the service port number of the Amazon RDS for PostgreSQL database. The default port number is 5432 .		
	Dat abase Name	Enter the name of the source Amazon RDS for PostgreSQL database.		
	Dat abase Account	Enter the account for the Amazon RDS for PostgreSQL database. For more information about permissions required for the account, see Permissions required for database accounts.		
		Enter the password for 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 Instance.		
	Instance Region	Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.		
	RDS Instance ID	Select the ID of the destination ApsaraDB RDS for PostgreSQL instance.		

Section	Parameter	Description
	Dat abase Name	Enter the name of the destination ApsaraDB RDS for PostgreSQL database. The name can be different from the name of the source Amazon RDS for PostgreSQL database.
Destination Database		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	Enter the database account of the destination ApsaraDB RDS for PostgreSQL instance. For more information about permissions required for the account, see Permissions required for database accounts.
		Enter the password for 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 The CIDR blocks of DTS servers are automatically added 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 objects to be migrated.



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



- 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.
- \circ If the task fails to pass the precheck, you can click the \bigcirc 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.
- 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.



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

8.5. Migrate data from an Amazon Aurora MySQL database to an ApsaraDB RDS for MySQL database

This topic describes how to migrate data from an Amazon Aurora MySQL database 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 configuring a data migration task, you can select all of the supported migration types to ensure service continuity.

Prerequisites

- The Public accessibility option of Amazon Aurora MySQL is set to **Yes**. This ensures that DTS can access Amazon Aurora MySQL over the Internet.
- The available storage space of the ApsaraDB RDS for MySQL instance is larger than the total space of the data in the Amazon Aurora MySQL instance.

> Document Version: 20220617

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
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 MySQL instance. DTS supports schema migration for the following types of objects: table, view, trigger, synonym, 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. Before a user can call views, stored procedures, and functions of the destination database, you must grant the read/write permissions to the user.

• Full dat a migration

DTS migrates historical data of the required objects from the source Amazon Aurora MySQL database to the destination ApsaraDB RDS for MySQL database.

Note During full data migration, concurrent INSERT operations cause segments 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 source Amazon Aurora MySQL database. Then, DTS synchronizes incremental data from the source Amazon Aurora MySQL database to the destination ApsaraDB RDS for MySQL database. Incremental data migration helps you 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 permission to perform SELECT operations on the objects to be migrated	The permission to perform SELECT operations on the objects to be migrated	The REPLICATION SLAVE permission, REPLICATION CLIENT permission, and permission to perform SELECT operations on the objects to be migrated
ApsaraDB RDS for MySQL	The read/write permissions for the objects to be migrated	The read/write permissions for the objects to be migrated	The read/write permissions for the objects to be migrated

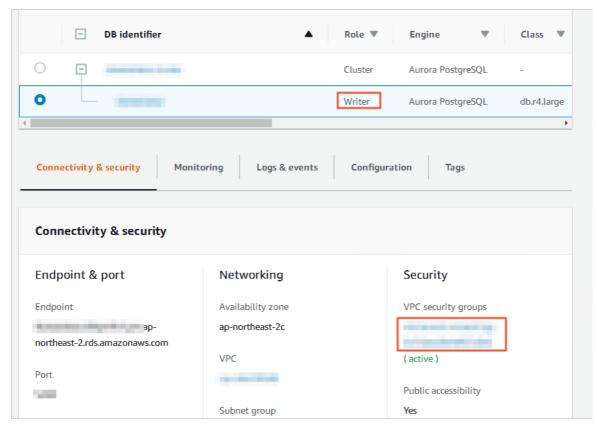
For more information about how to create and authorize a database account, see the following topics:

- Create an account for a user-created MySQL database and configure binary logging for an Amazon Aurora MySQL database
- Create an account for an RDS for MySQL instance and Change the permissions of an account for an RDS for MySQL instance for an ApsaraDB RDS for MySQL database

Preparations before data migration

- 1. Log on to the Amazon Aurora console.
- 2. Go to the **Basic information** page of the source Amazon Aurora MySQL instance.

- 3. Select the node that is set to the role of Writer.
- 4. In the **Connectivity & security** section, click the name of the VPC security group that corresponds to the Writer node.



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.

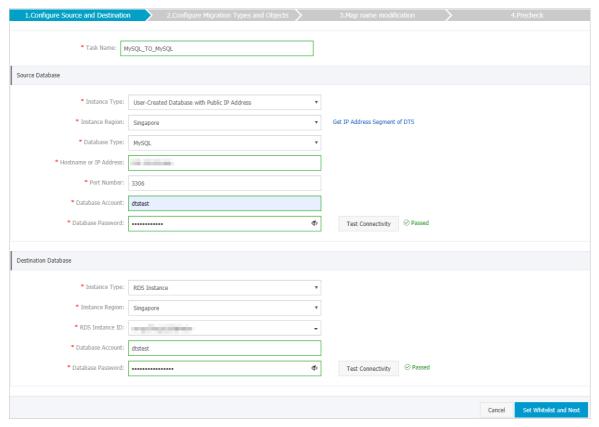
call mysql.rds_set_configuration('binlog retention hours', 24);



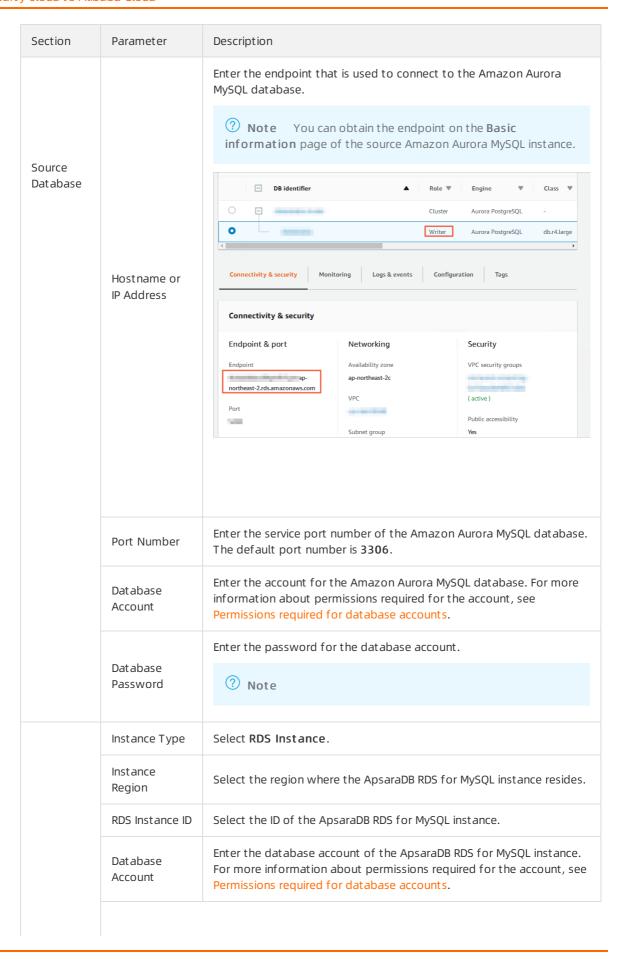
- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of Amazon Aurora MySQL must be enabled and the value of the binlog_format parameter must be set to row. If the version of MySQL is 5.6 or later, the value of the binlog_row_image parameter must be set to full.

Procedure

- 1.
- 2.
- 3.
- 4.
- 5. Configure the information about the source and destination databases for the data migration task.

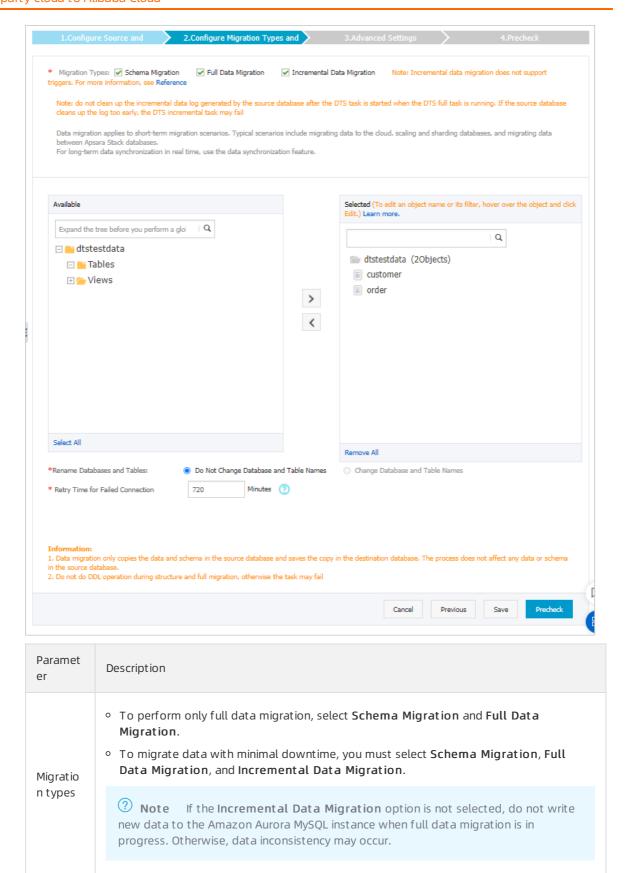


Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you use an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select User-Created Database with Public IP Address.
	Instance Region	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.
	Database Type	Select MySQL.



Destio atio	Parameter	Description	
n Database		Enter the password for the database account.	
	Dat abase Password	? Note	
Encryption		Select Non-encrypted or SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the RDS instance before configuring the data migration task. For more information, see Configure SSL encryption for an RDS for MySQL instance.	
Encryption	Note Encryption is available only in mainland China and Hong Kong(China).		

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



Paramet er	Description
	In the Available section, select the objects to be migrated and click the icon to add the objects to the Selected section.
Objects to be migrate d	 Note You can select databases, tables, or columns as the objects to be migrated. After an object is migrated to the destination instance, the name of the object remains unchanged. If you want an object to have a different name after the object is migrated to the ApsaraDB RDS for MySQL instance, you can use the object name mapping feature provided by DTS. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.

9.

10.

11.

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

8.6. Migrate data from an Amazon Aurora MySQL database to a PolarDB for MySQL database

This topic describes how to migrate data from an Amazon Aurora MySQL database to a PolarDB for MySQL database by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. When configuring 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 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 instance is larger than the total space of the data in the Amazon Aurora MySQL instance.

Notes

• 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 the source database does not have primary keys or UNIQUE constraints and fields are not required to be unique, duplicate data may exist in the destination database.
- DTS uses the ROUND (COLUMN, PRECISION) function to retrieve values from columns of the float or double data type. If the precision is not specified, 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 instance before configuring a data migration task.

Note For more information about how to create a database and the database naming conventions, see Create a database.

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

Billing

Migration type	Migration channel fee	Public network traffic fee
Schema migration or full data migration	Free of charge	Migrating data from Alibaba Cloud over the Internet
Incremental data migration	Billed. For more information, see Pricing.	incurs fees. For more information, see Pricing.

Migration types

Schema migration

DTS migrates the schemas of the required objects to the PolarDB 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.



- During schema migration, DTS changes the value of the SECURITY attribute in views, stored procedures, and functions from DEFINER to INVOKER.
- DTS does not migrate user information. Before a user can call views, stored procedures, and functions of the destination database, you must grant the read/write permissions to the user.
- Full data migration
 DTS migrates historical data of the required objects from the Amazon Aurora MySQL database to the destination PolarDB for MySQL database.

Note During full data migration, concurrent INSERT operations cause segments 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 source Amazon Aurora MySQL database. Then, DTS synchronizes incremental data from the source Amazon Aurora MySQL database to the destination PolarDB for MySQL database. Incremental data migration helps you ensure service continuity when you migrate data between MySQL databases.

Permissions required for database accounts

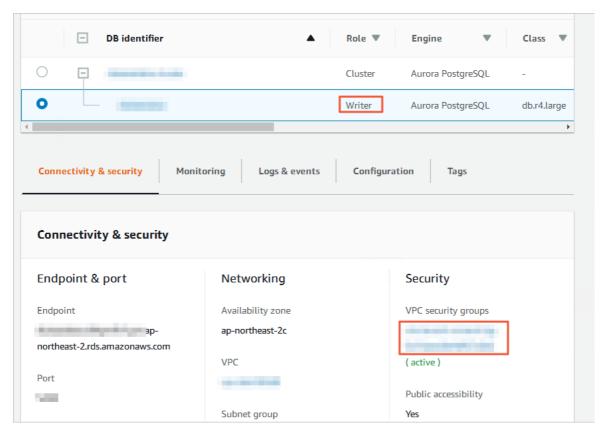
Database	Schema migration	Full data migration	Incremental data migration
Amazon Aurora MySQL	The permission to perform SELECT operations on the objects to be migrated	The permission to perform SELECT operations on the objects to be migrated	The REPLICATION SLAVE permission, REPLICATION CLIENT permission, SHOW VIEW permission, and permission to perform SELECT operations on the objects to be migrated
PolarDB for MySQL	The read/write permissions for the objects to be migrated	The read/write permissions for the objects to be migrated	The read/write permissions for the objects to be migrated

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

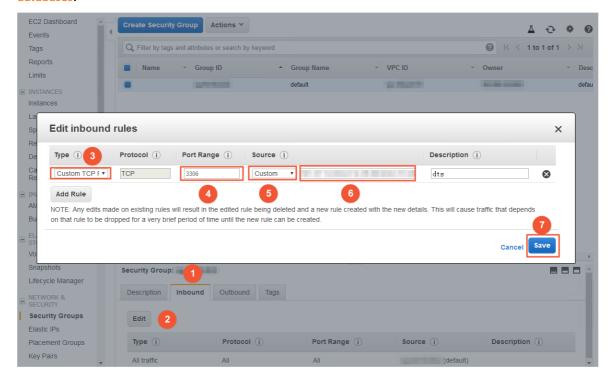
- Create an account for a user-created MySQL database and configure binary logging for an Amazon Aurora MySQL database
- Create database accounts for a PolarDB for MySQL database

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.



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.





- 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.

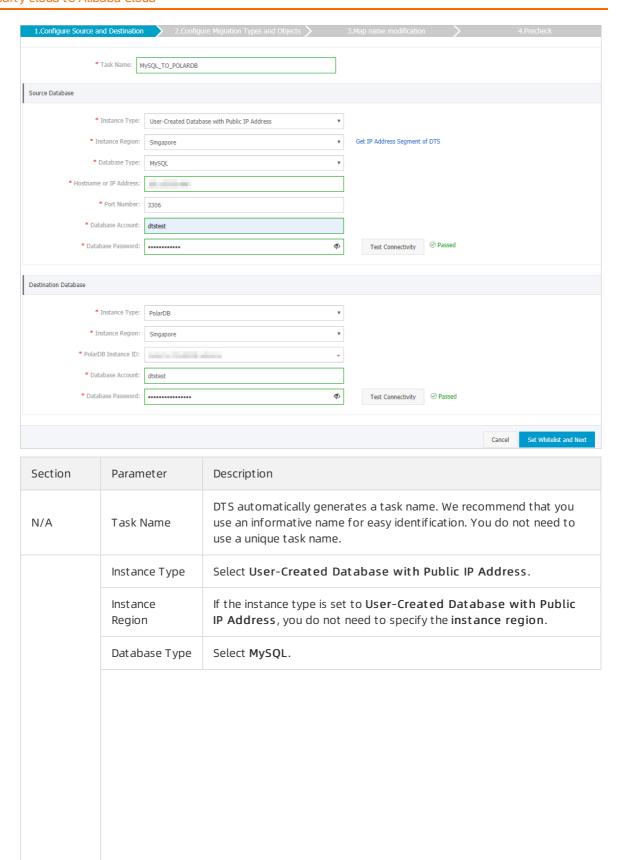
call mysql.rds_set_configuration('binlog retention hours', 24);

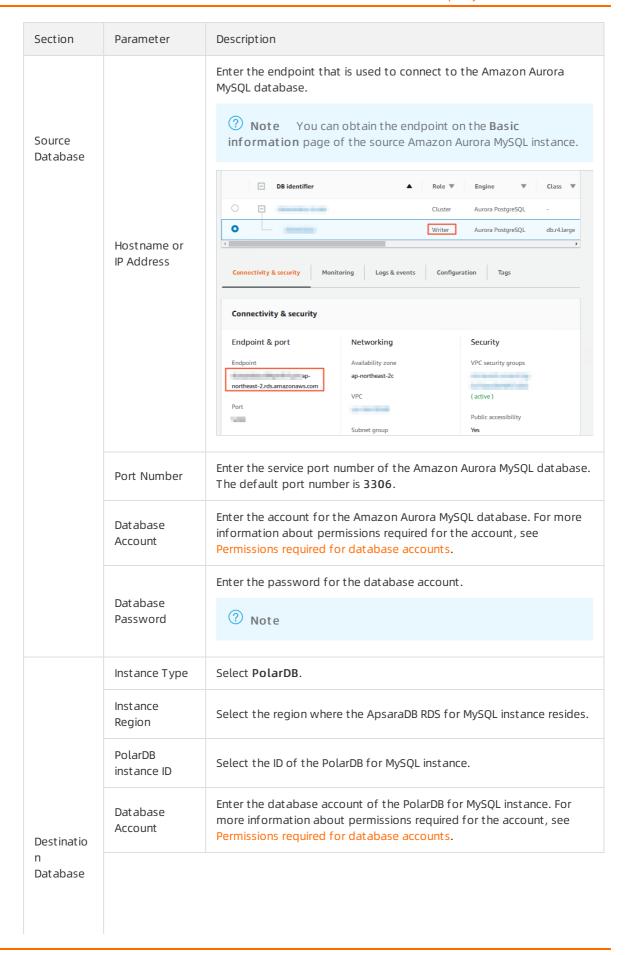


- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of the Amazon Aurora MySQL cluster must be enabled and the value of the binlog_format parameter must be set to row. If the MySQL version is 5.6 or later, the value of the binlog_row_image parameter must be set to full.

Procedure

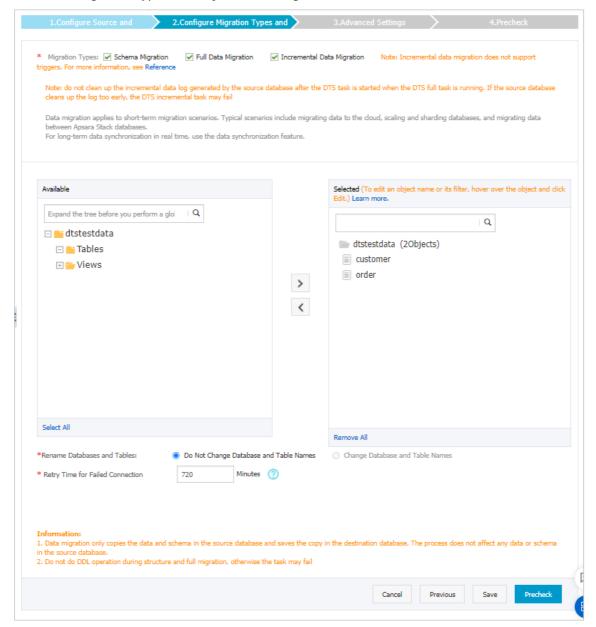
- 1.
- 2.
- 3.
- 4.
- 5. Configure the information about the source and destination databases for the data migration task.





Section	Parameter	Description
	Dat abase Password	Enter the password for the database account. Note

- 6. In the lower-right corner of the page, click Set Whitelist and Next.
 - Note The CIDR blocks of DTS servers are automatically added to the whitelist of the PolarDB for MySQL instance. This ensures that DTS servers can connect to the PolarDB for MySQL instance.
- 7. Select the migration types and objects to be migrated.



Paramet er	Description
Migratio n types	 To perform only full data migration, select Schema Migration and Full Data Migration. To migrate data with minimal downtime, you must select Schema Migration, Full Data Migration, and Incremental Data Migration. Note If the Incremental Data Migration option is not selected, do not write new data to the Amazon Aurora MySQL instance when full data migration is in progress. Otherwise, data inconsistency may occur.
	In the Available section, select the objects to be migrated and click the icon to add the objects to the Selected section.
Objects to be migrate d	 Note You can select databases, tables, or columns as the objects to be migrated. After an object is migrated to the destination instance, the name of the object remains unchanged. If you want an object to have a different name after the object is migrated to the PolarDB for MySQL instance, you can use the object name mapping feature provided by DTS. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.

9.

10. 11.

12. Switch your workloads to the PolarDB for MySQL instance.

8.7. Migrate full data from an Amazon Aurora PostgreSQL database to an ApsaraDB RDS for PostgreSQL database

This topic describes how to migrate full data from an Amazon Aurora PostgreSQL database to an ApsaraDB RDS for PostgreSQL database by using Data Transmission Service (DTS).

Prerequisites

• The version of the Amazon Aurora PostgreSQL database is 9.4, 9.5, 9.6, or 10.0.

- The **Public accessibility** option of Amazon Aurora PostgreSQL is set to **Yes**. This ensures that DTS can access Amazon Aurora PostgreSQL over the Internet.
- An ApsaraDB RDS for PostgreSQL instance is created. For more information, see Create an ApsaraDB RDS for PostgreSQL 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 space of the data in the Amazon Aurora PostgreSQL instance.

Notes

- DTS uses read and write resources of the source and destination databases during full data migration. This may increase the database load. 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 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 usage of the source and destination databases is less than 30%.
- DTS does not support incremental data migration. When using DTS for full data migration, you need to stop business operations that are related to the Amazon Aurora PostgreSQL database before migration. To ensure data consistency, do not write new data to the Amazon Aurora PostgreSQL database during 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 for pg_catalog	The permission to perform SELECT operations on the objects to be migrated
ApsaraDB RDS for PostgreSQL database	The create and usage permissions on the objects to be migrated	The owner permission for schemas

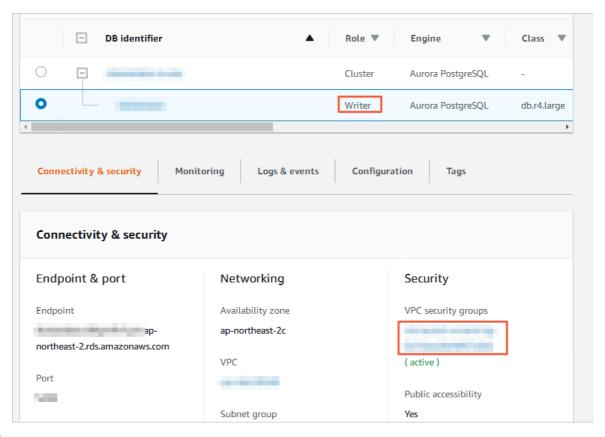
Process of full data migration

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: Modify Amazon inbound rules

- 1. Log on to the Amazon Aurora console.
- 2. Go to the Basic Information page of the source 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.



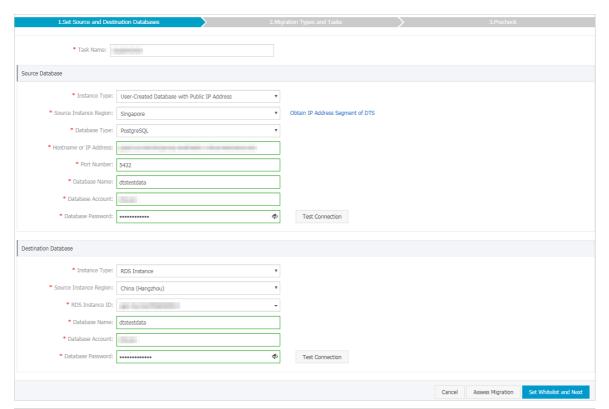
5.

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 destination database must be the same as that of the source database. For more information, see Create a database and Schema management.

Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Dat a 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 information about the source and destination databases for the data migration task.



Section	Parameter	Description	
N/A	Task Name	DTS automatically generates a task name. We recommend that you use 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 User-Created Database with Public IP Address , you do not need to specify the instance region.	
	Dat abase Type	Select PostgreSQL.	

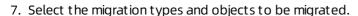
Section	Parameter	Description	
Source Dat abase	Hostname or IP Address	Enter the endpoint that is used to connect to the Amazon Aurora PostgreSQL database. ② Note You can obtain the endpoint on the Basic information page of the source Amazon Aurora PostgreSQL instance. Related ② Filter databases ③ DB identifier	
	Port Number Database	Enter the service port number of the Amazon Aurora PostgreSQL database. The default port number is 5432. Enter the name of the source Amazon Aurora PostgreSQL database.	
	Dat abase Account	Enter the account for the Amazon Aurora PostgreSQL database. For more information about permissions required for the account, see Permissions required for database accounts.	
	Dat abase Password	Enter the password for 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 RDS Instance.		Select RDS Instance.	
	Instance Region	Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.	
	RDS Instance ID	Select the ID of the destination ApsaraDB RDS for PostgreSQL instance.	

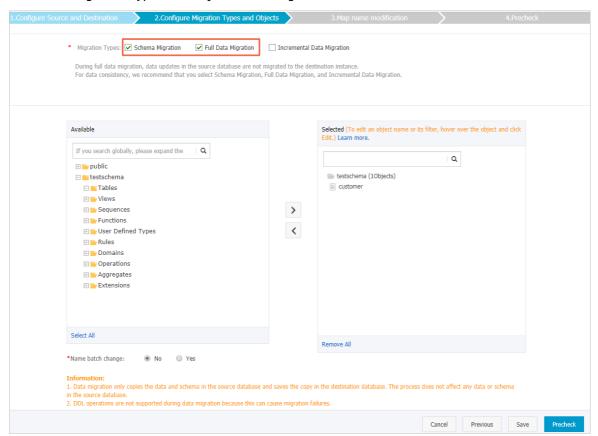
Section	Parameter	Description	
		Enter the name of the destination ApsaraDB RDS for PostgreSQL database. The name can be different from the name of the source Amazon Aurora PostgreSQL database.	
Destination Database	Dat abase Name	Note Before you configure the data migration task, you must create a database 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 more information about permissions required for the account, see Permissions required for database accounts.	
		Enter the password for 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 \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

Warning

- o 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.
- o 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.





Parameter	Description		
	Select Schema Migration and Full Data Migration . In this data migration task, incremental data migration is not supported.		
Migration types	Note To ensure data consistency, do not write new data into the Amazon Aurora PostgreSQL database during data migration.		
	In the Available section, select the objects to be migrated and click the icon to		
	add the objects to the Selected section.		
	? Note		
	 You can select databases, tables, or columns as the objects to be migrated. 		
Objects to be migrated	 After an object is migrated to the destination database, the name of the object remains unchanged. If you want an object to have a different name after the object is migrated to the ApsaraDB RDS for PostgreSQL instance, you can use the object name mapping feature provided by DTS. For more information about how to use this feature, see Object name mapping. 		
	 If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated. 		

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



- 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.
- $\circ \ \ \text{If the task fails to pass the precheck, you can click the} \\ \boxed{\textit{\textit{\textit{los}}}} \ \text{icon next to each failed item}$

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.
- 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.



 $12. \ \ Switch your workloads to the ApsaraDB RDS for PostgreSQL instance.$

8.8. Migrate full data from an Amazon RDS for SQL Server database to an ApsaraDB RDS for SQL Server database

This topic describes how to migrate full data from an Amazon RDS for SQL Server database to an ApsaraDB RDS for SQL Server database by using Data Transmission Service (DTS).

Prerequisites

- The **Public accessibility** option of Amazon RDS for SQL Server is set to **Yes**. This ensures that DTS can access Amazon RDS for SQL Server over the Internet.
- The database version of Amazon RDS for SQL Server is 2005, 2008, 2008 R2, 2012, 2014, or 2016.
- An ApsaraDB RDS for SQL Server instance is created. For more information, see Create an RDS for SQL Server instance.
- The available storage space of the destination ApsaraDB RDS for SQL Server database is larger than the total space of the data in the Amazon RDS for SQL Server database.

Notes

- 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 server-level role is not available 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 configuring a data migration task.

- **Note** For more information about how to create a database and the database naming conventions, see Create databases and accounts.
- DTS automatically resumes a failed data migration task. Before switching your workloads to the destination instance, you must stop or release the data migration task. Otherwise, the data from 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 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, and default.
- Full data migration
 DTS migrates historical data of the required objects from the source Amazon RDS for SQL Server database to the destination ApsaraDB RDS for SQL Server database.
 - Note Data of the sql_variant type cannot be migrated.

Billing

Migration type	Instance configurations	Internet traffic
Schema migration or full data migration	Free of charge	Migrating data from Alibaba Cloud over the Internet incurs fees. For more information, see Pricing.

Permissions required for database accounts

Database	Schema migration	Full data migration
Amazon RDS for SQL Server database	The permission to perform SELECT operations	The permission to perform SELECT operations
ApsaraDB RDS for SQL Server database	The read/write permissions	The read/write permissions

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

- CREATE USER for an Amazon RDS for SQL Server database
- Create an account for an RDS for SQL Server instance for an ApsaraDB RDS for SQL Server database

Full data migration

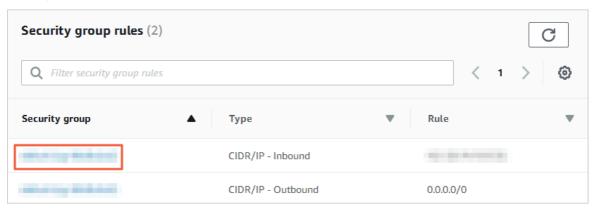
To avoid data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source SQL Server database in the following order:

- 1. Migrate the schemas of tables, views, synonyms, user-defined types, rules, defaults, and plan guides.
- 2. Perform full data migration.

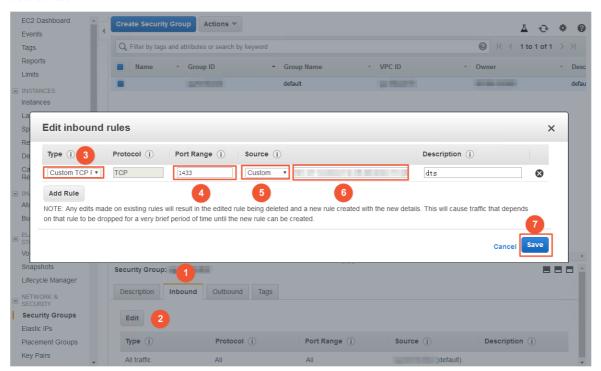
3. Migrate the schemas of SQL stored procedures, SQL functions, triggers, and foreign keys.

Preparations before data migration

- 1. Log on to the Amazon RDS Management Console.
- 2. Go to the Basic information page of the source 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.



4. On the **Security groups** page, click the Inbound tab in the Security group section. On the Inbound tab, click Edit to add 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.

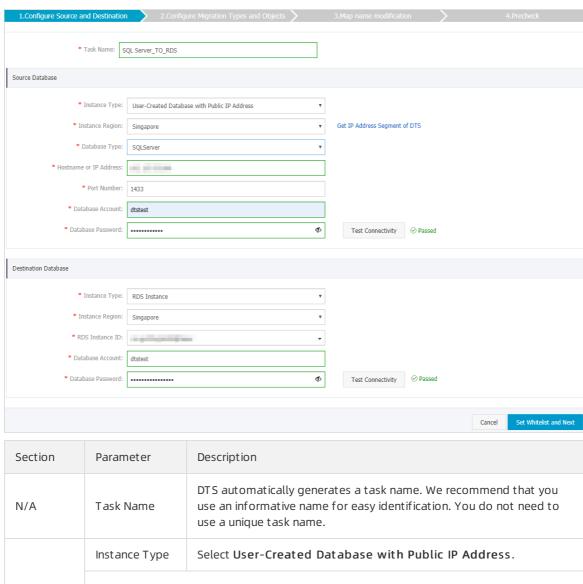


? Note

- You only need to add the CIDR blocks of DTS servers that are located in the same region
 as the destination database. For example, the source database is located in Singapore
 and the destination database is located in Hangzhou. You only need to add the CIDR
 blocks of DTS servers that are located in the China (Hangzhou) region.
- You can add all the required CIDR blocks to the inbound rule at one 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 information about the source and destination databases for the data migration task.



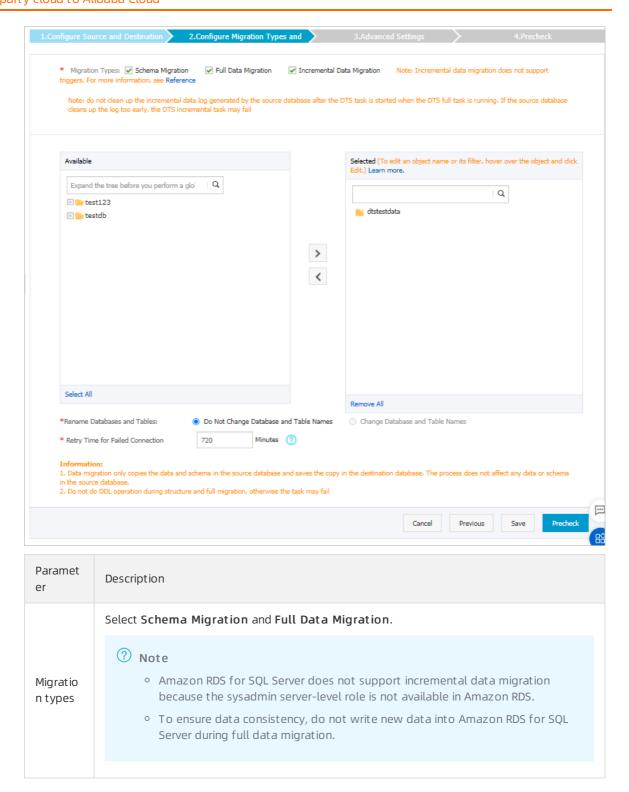
Section	Parameter	Description		
	Instance Region	If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.		
	Database Type	Select SQL Server .		
Source Dat abase	Hostname or IP Address	Enter the endpoint that is used to connect to the Amazon RDS for SQL Server database. Posterial Note You can obtain the endpoint on the Basic information page of the source Amazon RDS for SQL Server instance. Amazon RDS X		
	Port Number	Enter the service port number of the Amazon RDS for SQL Server database. The default port number is 1433.		
	Dat abase Account	Enter the account for the Amazon RDS for SQL Server database. For more information about permissions required for the account, see Permissions required for database accounts.		
	Dat abase Password	Enter the password for the database account. 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 RDS Instance.		
	Instance Region	Select the region where the ApsaraDB RDS for SQL Server instance resides.		
	RDS Instance ID	Select the ID of the ApsaraDB RDS for SQL Server instance.		

Section	Parameter	Description	
Destinatio n Database	Database Account	Enter the database account of the ApsaraDB RDS for SQL Server instance. For more information about permissions required for the account, see Permissions required for database accounts.	
		Enter the password for the database account.	
	Database 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 \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

Note The CIDR blocks of DTS servers are automatically added to the whitelist of the ApsaraDB for RDS instance. This ensures that DTS servers can connect to the ApsaraDB for RDS instance.

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



Paramet er	Description
	In the Available section, select the objects to be migrated and click the icon to add the objects to the Selected section.
	? Note
	 You can select databases, tables, or columns as the objects to be migrated.
Objects to be migrate d	 After an object is migrated to the destination instance, the name of the object remains unchanged. If you want an object to have a different name after the object is migrated to the ApsaraDB RDS for SQL Server instance, you can use the object name mapping feature provided by DTS. For more information about how to use this feature, see Object name mapping.
	 If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated.

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



- 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.
- 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.



12. Switch your workloads to the ApsaraDB RDS for SQL Server instance.

8.9. Migrate data from a MongoDB Atlas database to Alibaba Cloud

This topic describes how to migrate incremental data from a MongoDB Atlas database to Alibaba Cloud by using Data Transmission Service (DTS). DTS supports full data migration and incremental data migration. When you migrate data from a MongoDB Atlas database, you can select both migration types to ensure service continuity.

Prerequisites

The available storage space of the destination ApsaraDB for MongoDB instance is larger than the total size of the data in the MongoDB Atlas database.

Precautions

- During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the load of the database server. 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.
- You 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 MongoDB Atlas database and the destination ApsaraDB for MongoDB instance 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	Instance configuration	Internet traffic
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.

Permissions required for database accounts

Database Full data migration Incremental data migration

Database	Full data migration	Incremental data migration
MongoDB Atlas database	The read permission on the source database and the permission to perform the listDatabases operation	 The read permission on the source database, admin database, and local database The permission to perform the listDatabases operation
ApsaraDB for MongoDB instance	The read/write permissions on the destination database	The read/write permissions on the destination database

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

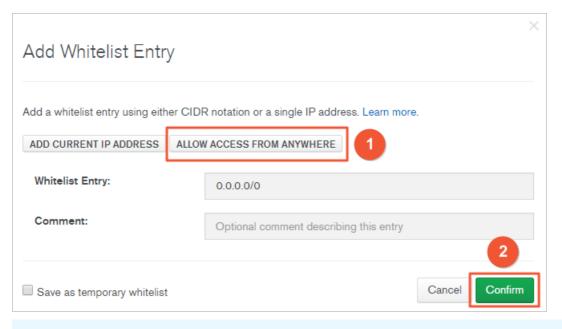
- Create User in MongoDB for a MongoDB Atlas database
- Manage MongoDB users through DMS for an ApsaraDB for MongoDB instance

Before you begin

- 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**.



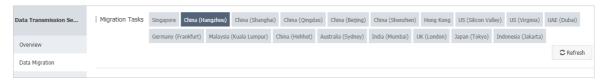
3. In the dialog box that appears, click ALLOW ACCESS FROM ANYWHERE, and click Confirm.



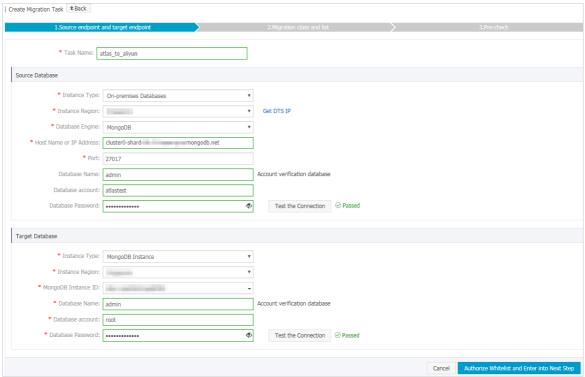
Note This step allows all IP addresses to access the MongoDB Atlas database. Delete this rule after data migration is complete.

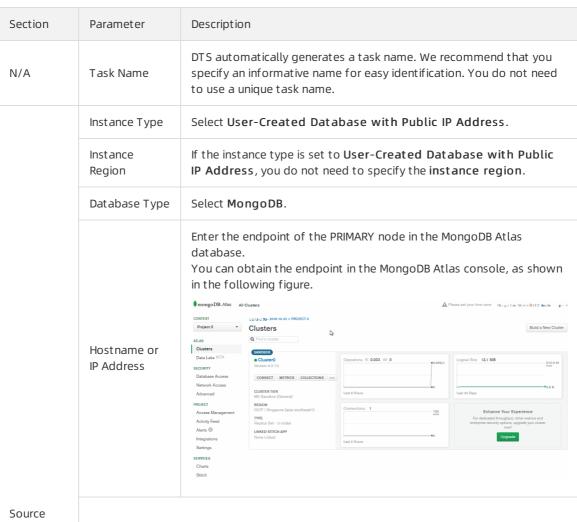
Procedure

- 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 for the data migration task.

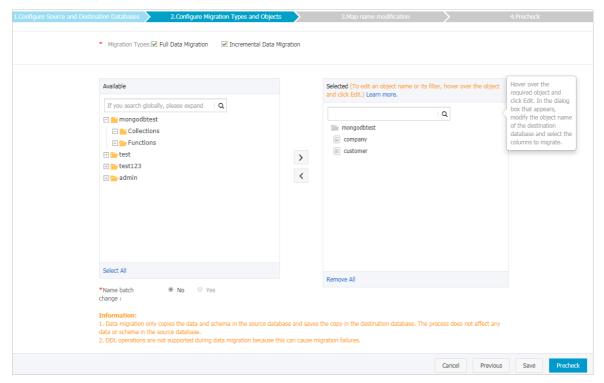




Database

Section	Parameter	Description	
	Port Number	Enter the service port number of the MongoDB Atlas database. The default port number is 27017 .	
		Enter the name of the authentication database. The database account is created in this database.	
	Dat abase Account	Enter the account of the MongoDB Atlas database. For more information about the permissions that are required for the account, see Permissions required for database accounts.	
	Dat abase Password	Phote 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.	
	Encryption	Select SSL-encrypted.	
	Instance Type	Select MongoDB Instance.	
	Instance Region	Select the region where the ApsaraDB for MongoDB instance resides.	
	MongoDB Instance ID	Select the ID of the ApsaraDB for MongoDB instance.	
	Database	Enter the name of the authentication database. The database account is created in this database.	
Destinatio	Name	Note If the database account is root, enter admin.	
n Database	Dat abase Account	Enter the database account of the ApsaraDB for MongoDB instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.	
		Enter the password of the destination 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 for MongoDB instance. This ensures that DTS servers can connect to the ApsaraDB for MongoDB instance.

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

Parameter	Description		
	To perform only full data migration, select only Full Data Migration.		
Migration Types	Note If you select only Full Data Migration, do not write data to the MongoDB Atlas database during data migration. This ensures data consistency between the source and destination databases.		
	 To ensure service continuity during data migration, select both Full Data Migration and Incremental Data Migration. 		

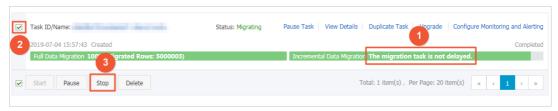
Parameter	Description
	 Select objects from the Available section and click the icon to move the objects to the Selected section. ? Note You cannot migrate data from the admin or local database.
Objects	 You can select databases, collections, or functions as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains unchanged. You can change the names of the objects that are migrated to the ApsaraDB for MongoDB instance by using the object name mapping feature. For more information about how to use this feature, see Object name mapping.

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



- 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.
- o If the task fails to pass the precheck, click the icon next to each failed item to view details. Troubleshoot the issues based on the causes 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 migration task.
 - Full data migration
 We recommend that you do not manually stop a data migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task automatically stops.
 - Increment al data migration
 An increment al data migration task does not automatically stop. You must manually stop the migration task.
 - ? Note Select an appropriate time to manually stop the migration task. For example, you can stop the migration 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 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. After the status of incremental data migration changes to The migration task is not delayed, manually stop the migration task.



12. Switch your workloads to the ApsaraDB for MongoDB instance.

9.Migrate using the same Alibaba Cloud account

9.1. 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 source and destination RDS instances use the same database engines.

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 data consistency, 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.

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	
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.
- Full dat a migration

DTS migrates historical data of the required objects from the source RDS instance to the destination RDS instance.

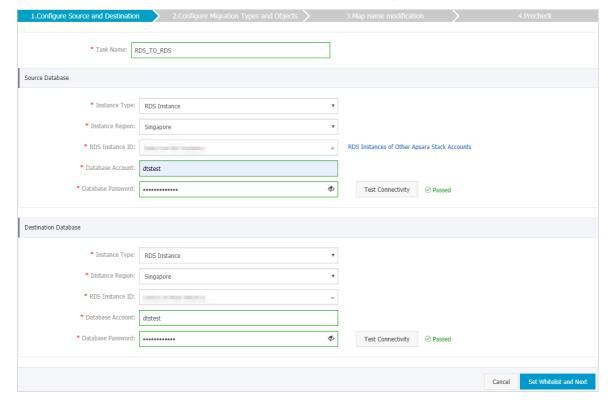
Incremental data migration
 After full data migration is complete, DTS synchronizes incremental data from the source RDS instance to the destination RDS instance. Incremental data migration allows you to ensure service continuity when you migrate data between RDS instances.

Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Source RDS instance	The read/write permissions	The read/write permissions	The read/write permissions
Destination RDS instance	The read/write permissions	The read/write permissions	The read/write 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 for the data migration task.



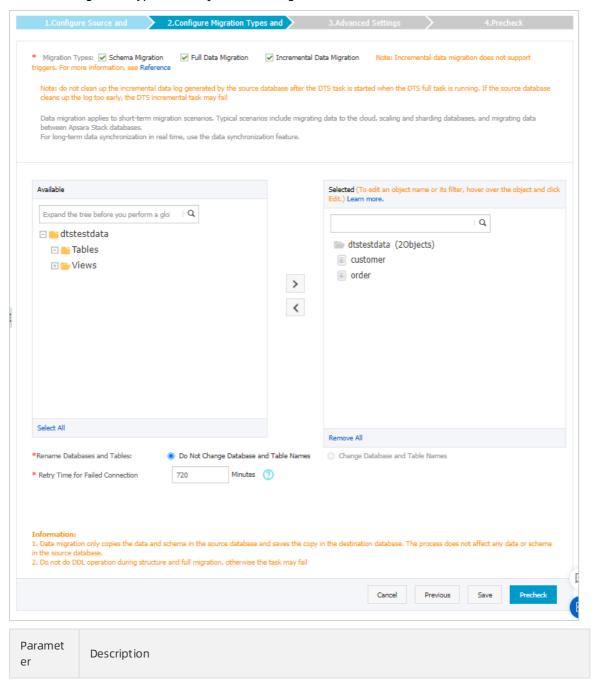
Section	Parameter	Description	
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.	
	Instance Type	Select RDS Instance.	
	Instance Region	Select the region where the source RDS instance resides.	
	RDS Instance ID	Select the ID of the source RDS instance. Note 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.	
	Dat abase Name	Enter the name of the source database in the ApsaraDB RDS for PostgreSQL instance. Note This parameter is required only if the database engine of the RDS instance is PostgreSQL.	
	Dat abase Account	Enter the database account of the source RDS instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.	
Source Dat abase	Dat abase Password	Pote After you specify the information about the selfmanaged 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.	
	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 for an RDS MySQL instance. 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 mainland China and the China (Hong Kong) region.	

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	Note 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.		
	Database	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.		
	Name	Note This parameter is required only if the database engine of the RDS instance is PostgreSQL.		
Destinatio n Database	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 destination 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.		
	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 for an RDS MySQL instance.		
		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 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 and destination RDS instances. This ensures that DTS servers can connect to the source and destination RDS instances.

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

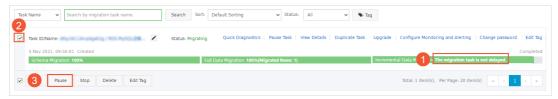


Paramet er	Description
Migratio n Types	 Select the migration types based on your business requirements. The migration types must be supported by the database engine. To perform only full data migration, select Schema Migration and Full Data Migration. To ensure service continuity during data migration, select Schema Migration, Full Data Migration, and Incremental Data Migration. Note If Incremental Data Migration is not selected, do not write data to the source RDS instance during full data migration. This ensures data consistency between the source and destination instances.
Available	Select objects from the Available section and click the icon to move the objects to the Selected section. Note 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. 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. 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.

8.

- 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.
 - o Full dat a 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 self-managed 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.



9.2. 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 migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance, you can select 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.

? Note The available storage space of the ApsaraDB RDS for MySQL instance must be 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.

• DTS automatically resumes a failed data migration 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
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 object: 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.
- o DTS does not migrate user information. Before a user can call views, stored procedures, and functions of the destination instance, you must grant the read/write permissions to the user.
- Full dat a migration DTS migrates historical data of the required objects to the destination instance.



? Note

- o During full data migration, concurrent INSERT operations cause segments in the tables of the destination instance. After full data migration is completed, the tablespace of the destination instance is larger than that of the source instance.
- o To ensure successful data migration, do not perform DDL operations on the source instance during full data migration.
- Incremental data migration After full migration is completed, DTS retrieves binary log files from the source instance, and migrates incremental data to the destination instance in real time.

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	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

Permissions required for database accounts

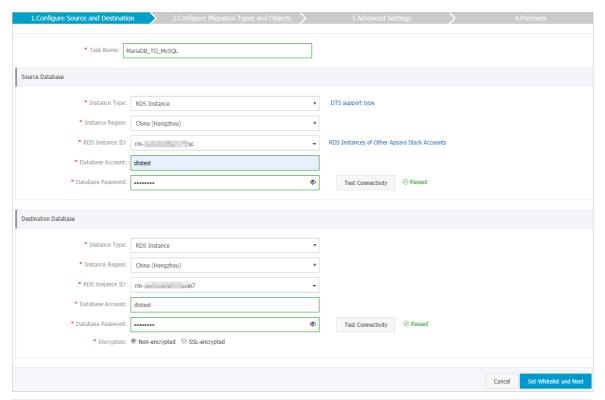
Database	Schema migration	Full data migration	Incremental data migration
Source ApsaraDB RDS for MariaDB TX instance	The SELECT permission	The SELECT permission	The REPLICATION SLAVE permission, the REPLICATION CLIENT permission, the SHOW VIEW permission, and the permission to perform SELECT operations on the required objects
Destination ApsaraDB RDS for MySQL instance	The read/write permissions	The read/write permissions	The read/write permissions

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

- Create accounts and databases for an ApsaraDB RDS for MariaDB TX instance
- Create an account for an ApsaraDB RDS for MySQL instance and Change the permissions of an account for an RDS 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 for the data migration task.



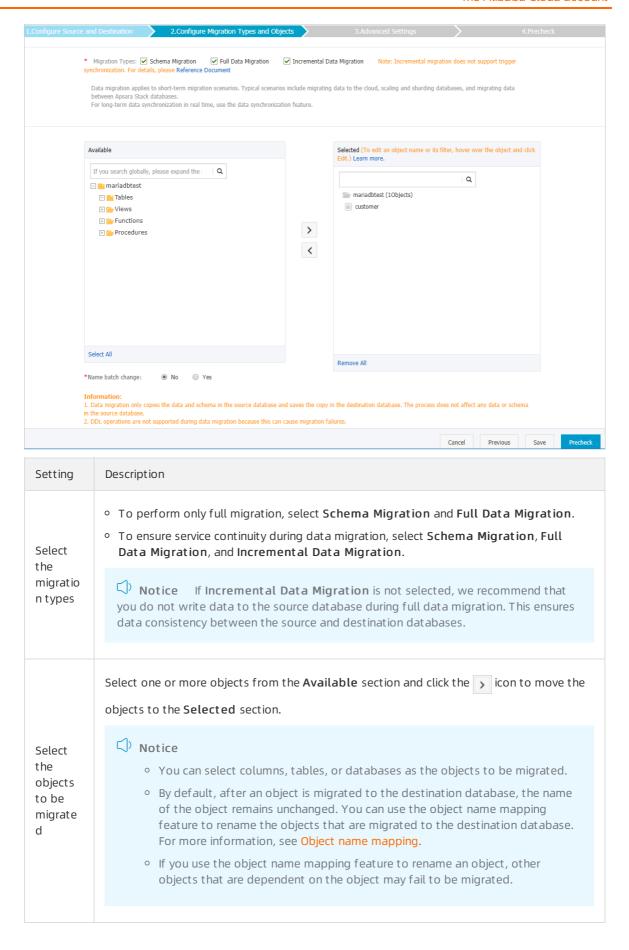
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you use an informative name for easy identification. You do not need to use a unique task name.
Source Dat abase	Instance Type	Select RDS Instance.
	Instance Region	Select the region where the source ApsaraDB RDS for MariaDB TX instance resides.
	RDS Instance ID	Select the ID of the source ApsaraDB RDS for MariaDB TX instance.
	Dat abase Account	Enter 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	Enter the password of the database account. After you specify the source database parameters, click Test Connectivity next to the Database Password parameter to verify whether the parameters are valid.
		Note If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the parameters based on the check results.

Section	Parameter	Description
Destination Database	Instance Type	Select RDS Instance.
	Instance Region	Select the region where the destination ApsaraDB RDS for MySQL instance resides.
	RDS Instance ID	Select the ID of the destination ApsaraDB RDS for MySQL instance.
	Dat abase Account	Enter 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.
	Database Password	Enter the password of the database account. After you specify the destination database parameters, click Test Connectivity next to the Database Password parameter to verify whether the parameters are valid.
		Note If the specified parameters are valid, the Passed message appears. If the Failed message appears, click Check next to Failed. Modify the parameters based on the check results.
	Encryption	Select Non-encrypted or SSL-encrypted. If you select SSL-encrypted, you must have enabled SSL encryption for the ApsaraDB RDS for MySQL instance. For more information, see Configure SSL encryption for an RDS MySQL instance.
		? Note The Encryption parameter is available only for data migration instances that reside in mainland China and Hong Kong (China).

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

Note The CIDR blocks of DTS servers are automatically added to the whitelist of the destination instance. This ensures that DTS servers can connect to the destination instance.

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



Setting	Description
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	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.
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.
s online DDL operatio ns on	Note If online DDL operations generate a large amount of data, the data migration task may be delayed.
	 No: 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.
	Note If you select No, the tables in the destination database may be locked.
the source table	

8. In the lower-right corner of the page, click $\mbox{\bf Precheck}.$



- 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.
- \circ If the task fails to pass the precheck, you can click the \bigcirc icon next to each failed item

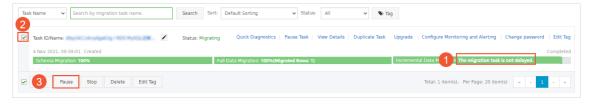
to view details.

- 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.
- 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.
- Increment al 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 the read/write permissions. After the data migration is completed, you must delete the database accounts of both the ApsaraDB RDS for MariaDB TX instance and the ApsaraDB RDS for MySQL instance to ensure database security.

9.3. Migrate data from an ApsaraDB RDS for MySQL instance to an Apsara PolarDB for MySQL cluster

Apsara PolarDB is a next-generation relational database service that is developed by Alibaba Cloud. The service 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 an Apsara PolarDB for MySQL cluster by using Data Transmission Service (DTS).

Prerequisites

The destination Apsara PolarDB for MySQL cluster is created. For more information, see Create a PolarDB 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 tablespace of the destination database is larger than that of the source database.
- 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.

Limits

• DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function.



• DTS does not migrate user information from the source database. Before a user can call views, stored procedures, and functions of the destination database, you must grant the read/write permissions to the user.

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 Alibaba Cloud over the Internet. For more
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

SQL operations that can be synchronized during incremental data migration

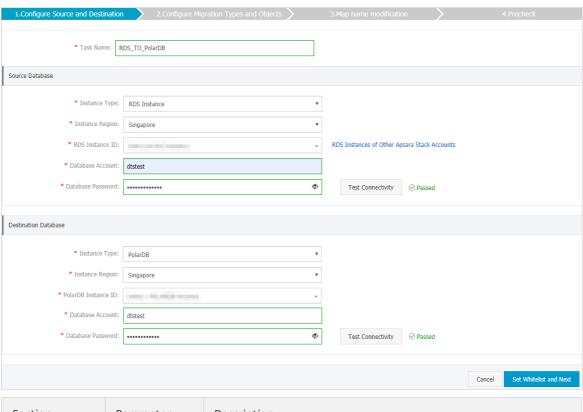
Operatio n type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

Permissions required for database accounts

Database	Required permission
ApsaraDB RDS for MySQL	The read permission on the objects to be migrated
Apsara PolarDB for MySQL	The read/write permissions on the objects to be migrated

Note For more information about how to create and authorize a database account, see Create an account for an ApsaraDB RDS for MySQL instance and Create database accounts.

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Dat a 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 for the data migration task.



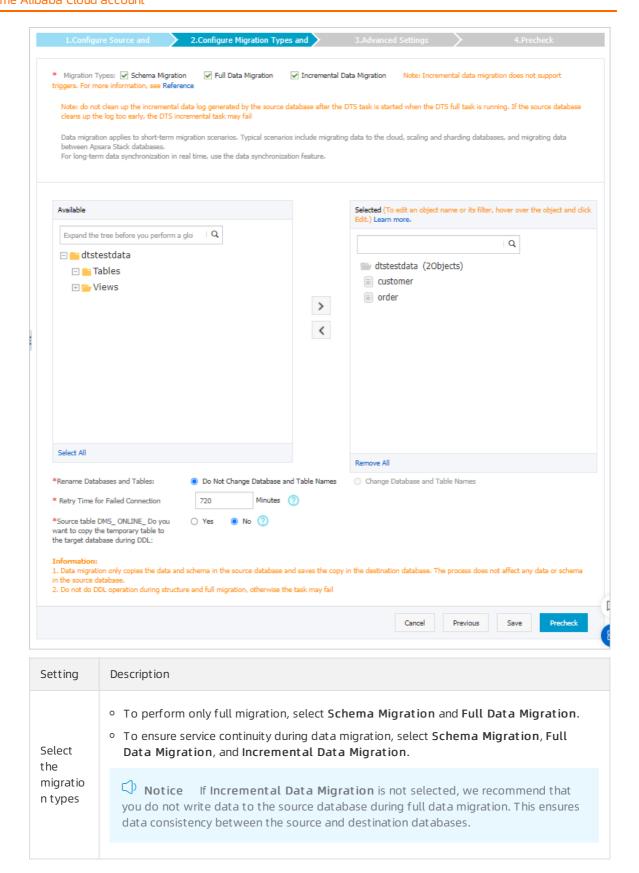
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select RDS Instance.
	Instance Region	Select the region where the source RDS instance resides.
	Dat abase Account	Enter the database account of the source 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 source database account.
Source Dat abase	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.

Section	Parameter	Description
	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 for an RDS MySQL instance. ② Note The Encryption parameter is available only for regions in mainland China and the China (Hong Kong) region.
	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. For more information about the permissions that are required for the account, see Permissions required for database accounts.
Database	Database Password	Enter the password of the destination 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.

6. In the lower-right corner of the page, click \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .

Note DTS adds the CIDR blocks of DTS servers to the whitelists of the source ApsaraDB RDS for MySQL instance and the destination Apsara PolarDB for MySQL cluster. This ensures that DTS servers can connect to the source instance and the destination cluster.

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



Setting	Description
Select the objects to be migrate d	Select one or more objects from the Available section and click the icon to move the objects to the Selected section. Notice You can select columns, tables, or databases as the objects to be migrated. 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. 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 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	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.

Setting	Description
Specify whether to copy tempora ry tables to the destinati on	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.
databas e when DMS perform s online DDL operatio ns on the source table	 No: 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.
	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**.



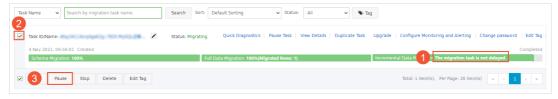
- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- \circ If the task fails to pass the precheck, you can click the \bigcirc icon next to each failed item

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.
- 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.

9.4. Migrate data between Apsara PolarDB for MySQL clusters

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

Note An Apsara PolarDB for MySQL cluster of an earlier version cannot be upgraded to 8.0. However, you can create an Apsara PolarDB for MySQL cluster of version 8.0, and then migrate data to this cluster. Before you migrate data between different versions of Apsara 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 Apsara PolarDB for MySQL clusters are created. For more information, see Create a PolarDB MySQL cluster.
- The binary logging feature is enabled for the source Apsara 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.
- 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.

Migration types

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

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

Billing

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

SQL operations that can be synchronized during incremental data migration

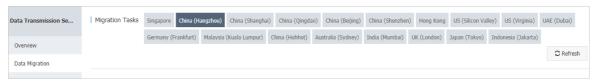
Operatio n type	SQL statement
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

Permissions required for database accounts

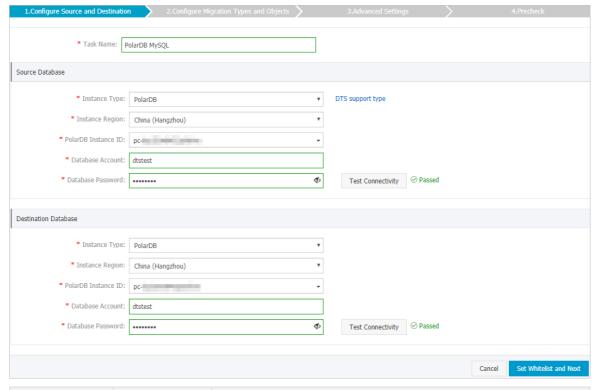
Database	Required permission
Source Apsara PolarDB for MySQL cluster	The read permission on the objects to be migrated
Destination Apsara PolarDB for MySQL cluster	The read/write permissions on the objects to be migrated

Note For more information about how to create and authorize a database account, see Create database accounts.

- 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 for the data migration task.



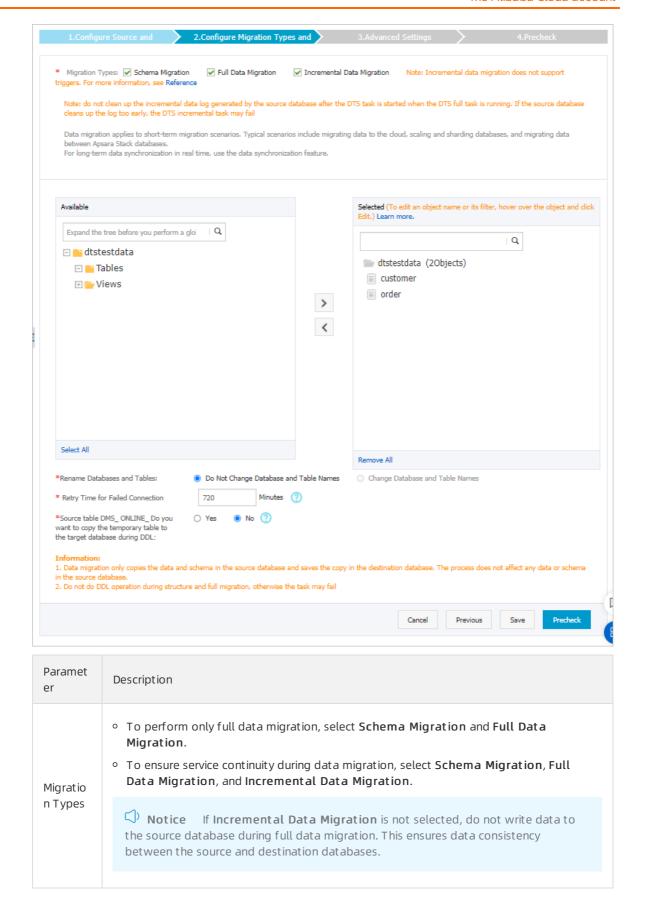
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select PolarDB.
	Instance Region	Select the region where the source PolarDB cluster resides.

Section	Parameter	Description
	PolarDB Instance ID	Select the ID of the source PolarDB cluster.
Source Dat abase	Dat abase Account	Enter the database account of the source PolarDB cluster.
		Enter the password of the source 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 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 destination 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 objects to be migrated.



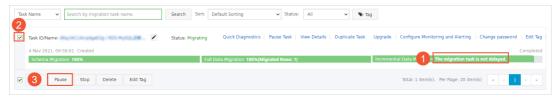
Paramet er	Description		
	Select objects from the Available section and click the icon to move the objects to the Selected section.		
Objects	 Notice You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains unchanged. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated. 		

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



- Before you can start the data migration task, a precheck is performed. A data migration task can be started only if it passes the precheck.
- o If the task fails to pass the precheck, click the icon next to each failed item to view details. Troubleshoot the issues based on the causes and perform 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 a data migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task automatically stops.
 - Schema migration, full data migration, and incremental data migration
 An incremental data migration task does not automatically stop. You must manually stop the migration task.
 - Notice Select an appropriate time to manually stop the migration task. For example, you can stop the migration task during off-peak hours or before you switch your workloads to the destination instance.
 - 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. After the status of incremental data migration changes to The migration task is not delayed, manually stop the migration task.



12. Switch your workloads to the destination PolarDB cluster.

What's next

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

9.5. Migrate data from an Apsara PolarDB for MySQL cluster to an ApsaraDB RDS for MySQL instance

ApsaraDB for RDS is a stable, reliable, and scalable online database service. ApsaraDB for 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 an Apsara 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 source Apsara PolarDB for MySQL cluster. For more information, see Enable-binlogging.
- 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.



• 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 Alibaba Cloud over the Internet. For more
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	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

Permissions required for database accounts

Database	Permission
Apsara PolarDB for MySQL	The read permission on the objects to be migrated

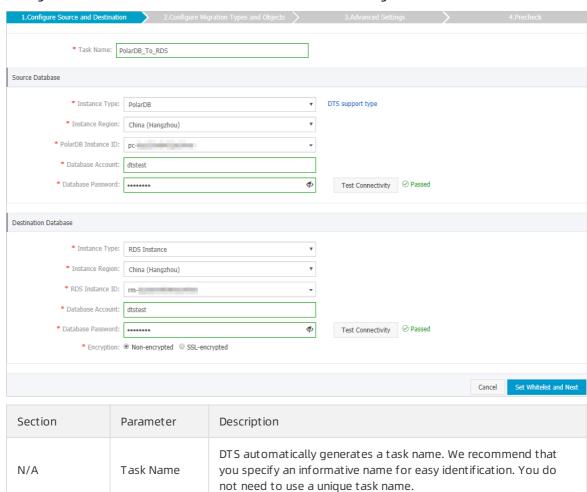
Database	Permission
ApsaraDB RDS for MySQL	The read/write permissions on the objects to be migrated

Note For more information about how to create and authorize a database account, see Create database accounts and Create an account for 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 RDS instance resides.



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



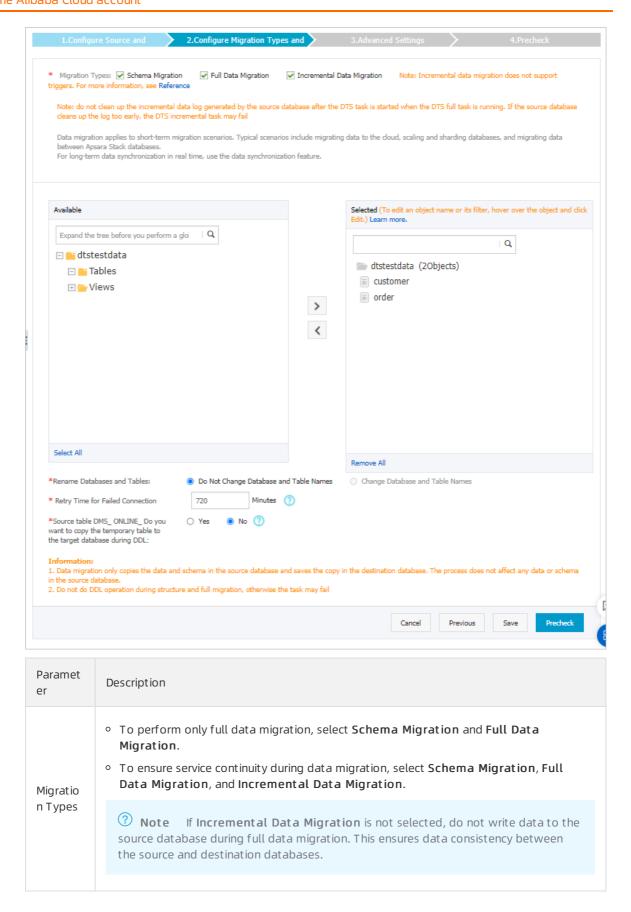
Section	Parameter	Description	
	Instance Type	Select PolarDB .	
	Instance Region	Select the region where the source PolarDB cluster resides.	
	PolarDB Instance ID	Select the ID of the source PolarDB cluster.	
Source	Dat abase Account	Enter the database account of the source PolarDB cluster. For more information about the permissions that are required for the account, see Permissions required for database accounts.	
Database		Enter the password of the source 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 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 destination 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.	

Section	Parameter	Description
Encryptic	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 for an RDS MySQL instance.
		Note 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.

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 objects to be migrated.



Paramet er	Description		
	Select objects from the Available section and click the icon to move the objects to the Selected section.		
Objects	 Note You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains unchanged. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated. 		

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



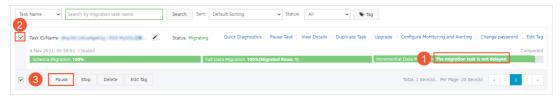
- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- \circ If the task fails to pass the precheck, you can click the or icon next to each failed item

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.
- 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 RDS instance.

9.6. Migrate data between PolarDB clusters compatible with Oracle

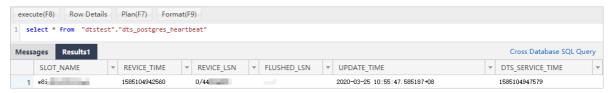
Apsara PolarDB is a next-generation relational database service developed by Alibaba Cloud. It is compatible with MySQL, PostgreSQL, and Oracle database engines. Apsara PolarDB provides superior performance in storage and computing to meet diverse requirements of enterprises. This topic describes how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

Prerequisites

The tables to be migrated 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%.
- 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 delay time of incremental data migration is accurate, DTS adds a heart beat table named dts_postgres_heartbeat to the source database. The following figure shows the schema of the heart beat table.



• DTS automatically resumes a failed data migration 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.

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
Incremental data migration	Charged. For more information, see Pricing.	information, see Pricing.

Migration types

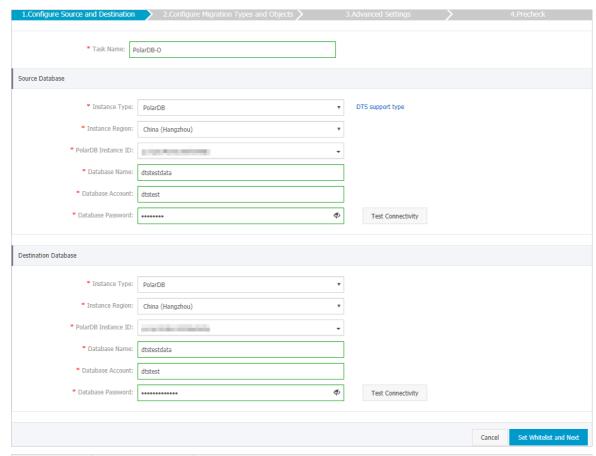
Migration type	Description
	DTS migrates the schemas of the required objects to the destination PolarDB cluster. DTS supports schema migration for the following types of objects: table, view, synonym, trigger, stored procedure, function, package, and user-defined type.
Schema migration	Notice However, 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 PolarDB cluster to the destination PolarDB cluster.
Full data migration	Notice 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.
	DTS retrieves redo log files from the source PolarDB cluster. Then, DTS synchronizes incremental data from the source PolarDB cluster to the destination PolarDB cluster. The following SQL operations can be synchronized: • INSERT, UPDATE, and DELETE operations • CREATE TABLE operations
Incremental data migration	Notice The CREATE TABLE operations used to create partition tables or tables that contain functions cannot be synchronized.
	ALTER TABLE, DROP TABLE, RENAME TABLE, CREATE INDEX, and ADD INDEX operations Incremental data migration allows you to onsure continuity when your
	Incremental data migration allows you to ensure service continuity when you migrate data between PolarDB clusters.

Preparation

Set the value of the wal_level parameter to *logical* for the source PolarDB cluster. This setting ensures that logical decoding is supported in write-ahead logging (WAL). For more information, see Configure cluster parameters.

• Warning If you change the value of the wal_level parameter, the source PolarDB cluster is restarted. Perform this operation with caution.

- 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.

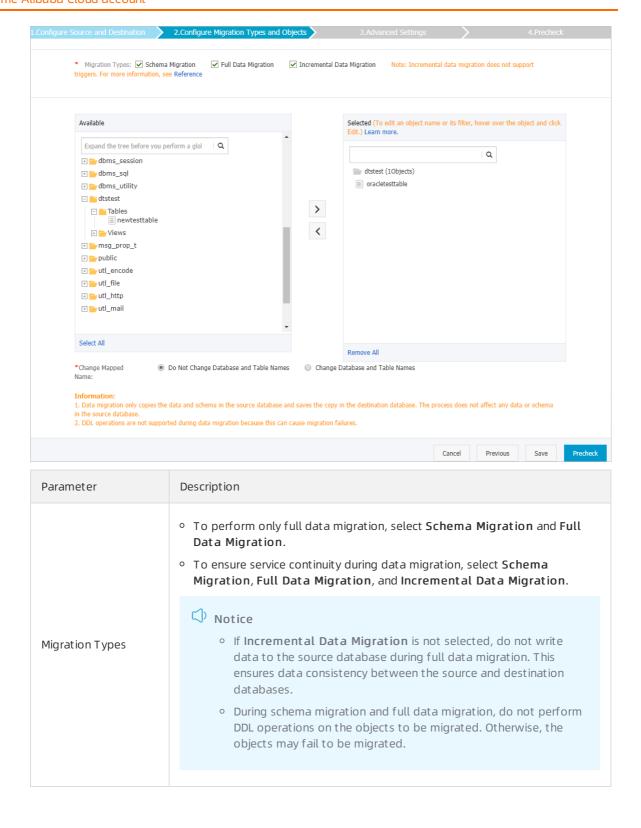


Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select PolarDB.
	Instance Region	Select the region where the source PolarDB cluster resides.
	PolarDB Instance ID	Select the ID of the source PolarDB cluster.

Section	Parameter	Description
Source	Dat abase Name	Enter the name of the source database.
Database	Dat abase Account	Enter the privileged account of the source PolarDB cluster. For more information, see Create database accounts.
		Enter the password of the source 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 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	Dat abase Account	Enter the database account of the destination PolarDB cluster. The account must have the owner permission on the database .
n Dat abase		Notice You can specify the database owner when you create a database.
	Dat abase Password	Enter the password of the destination database account.
		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**.

 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 objects to be migrated.



Parameter	Description	
	Select objects from the Available section and click the icon to move the objects to the Selected section. You can select columns, tables, or schemas as the objects to be migrated.	
Objects	 Notice After an object is migrated to the destination database, the name of the object remains unchanged. You can change the names of the objects that are migrated to the destination PolarDB cluster by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated. 	

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



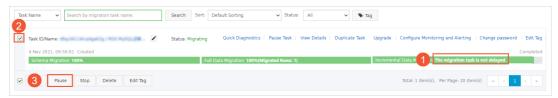
- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- $\circ \ \ \text{If the task fails to pass the precheck, you can click the} \\ \underline{\textit{\textit{\textit{los}}}} \ \text{icon next to each failed item}$

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.
- 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.



Related topics

Switch workloads to the destination database

9.7. Migrate data from a standalone MongoDB 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 designed to store non-core enterprise data, for example, in development and testing scenarios. Replica set instances and sharded cluster instances are more suitable for production scenarios. This topic describes how to migrate data from a standalone instance to a replica set instance or a 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

- During full data migration, DTS uses read and write resources of the source and destination
 databases. This may increase the load of the database server. 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 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	Instance configuration	Internet traffic
Full data migration	Free of charge	Free of charge

Migration types

Full data migration: DTS migrates all historical data of the required objects from the source MongoDB database to the destination MongoDB database.



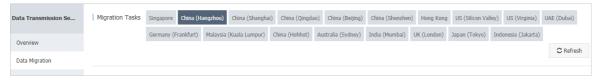
Note The following types of objects are supported: database, collection, and index.

Permissions required for database accounts

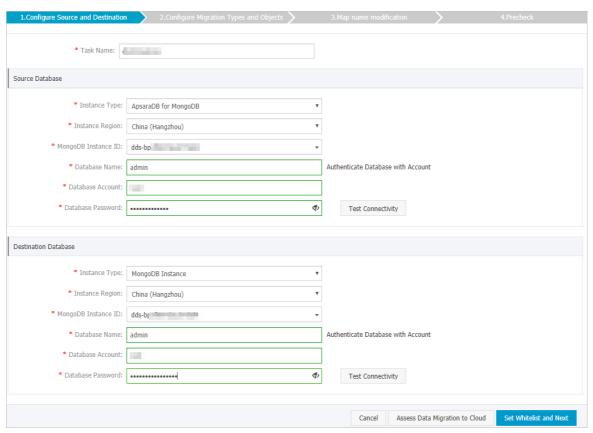
Instance	Permission
Source ApsaraDB for MongoDB instance	The read permission on the source database
Destination ApsaraDB for MongoDB instance	The read/write permissions on the destination database

For more information about how to create and authorize a database account, see Use DMS to manage MongoDB users.

- 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 for the data migration task.



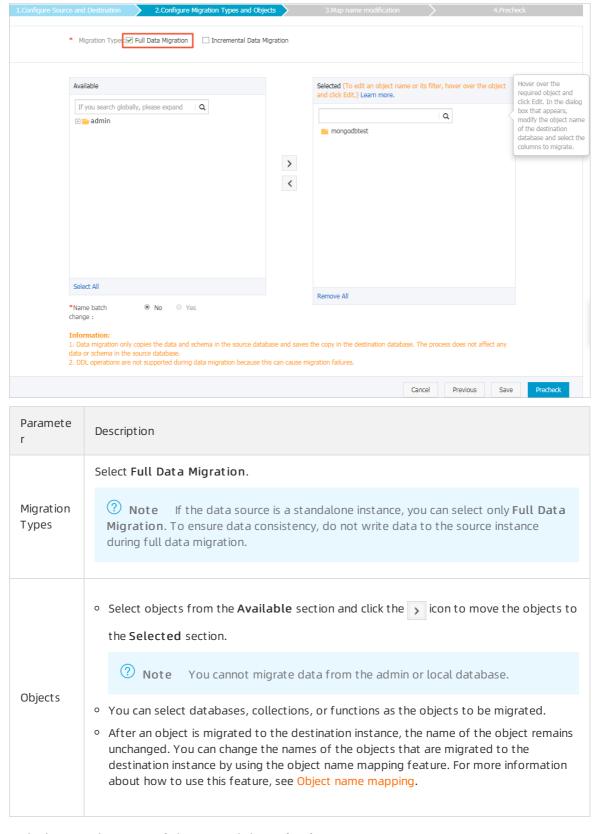
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select ApsaraDB for MongoDB.
	Instance Region	Select the region where the source ApsaraDB for MongoDB instance resides.
	MongoDB Instance ID	Select the ID of the source ApsaraDB for MongoDB instance.
	Dat abase Name	Enter the name of the authentication database. The database account is created in this database.
		? Note If the database account is root, enter admin.
Source	Dat abase Account	Enter the database account of the source ApsaraDB for MongoDB instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.
Database		

Section	Parameter	Description
	Dat abase Password	Enter the password of the source 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 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.
Destinatio n Database	Dat abase Name	Enter the name of the authentication database. The database account is created in this database. Note If the database account is root, enter admin.
	Dat abase Account	Enter the database account of the destination ApsaraDB for MongoDB 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 destination 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.

6. In the lower-right corner, 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.

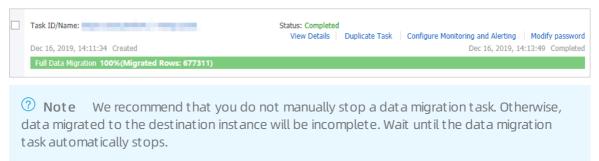
7. Select the migration types and objects to be migrated.



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



- 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.
- o If the task fails to pass the precheck, click the icon next to each failed item to view details. Troubleshoot the issues based on the causes 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 and select Data Transmission Service (Pay-As-You-Go) Service Terms.
- 11. Click **Buy and Start** to start the data migration task.



12. Switch your workloads to the destination ApsaraDB for MongoDB instance.

References

- Connect to a replica set instance
- Connect to a sharded cluster instance
- Configure sharding to maximize the performance of shards

9.8. Migrate data from a replica set MongoDB instance to a sharded cluster instance

This topic describes how to migrate data from a replica set MongoDB instance to a 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

• During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the load of the database server. 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	Instance configuration	Internet traffic
Full data migration	Free of charge.	Charged only when data is migrated from Alibaha
Incremental data migration	Charged. For more information, see DTS pricing.	Charged only when data is migrated from Alibaba Cloud over the Internet. For more information, see DTS pricing.

Migration types

Migration type	Description
Full data migration	DTS migrates all historical data of the required objects from the source MongoDB database to the destination MongoDB database.
	Note The following types of objects are supported: database, collection, and index.
Incremental data migration	After full data migration is complete, DTS synchronizes incremental data from the source MongoDB database to the destination MongoDB database.
	 Note The create and delete operations that are performed on databases, collections, and indexes can be synchronized. The create, delete, and update operations that are performed on documents can be synchronized.

Permissions required for database accounts

Instance	Full data migration	Incremental data migration
ApsaraDB for MongoDB replica set instance	The read permission on the source database	The read permission on the source database, admin database, and local database
ApsaraDB for MongoDB sharded cluster instance	The read/write permissions on the destination database	The read/write permissions on the destination database

Note For more information about how to create and authorize a database account, see Use DMS to manage MongoDB users.

Before you begin

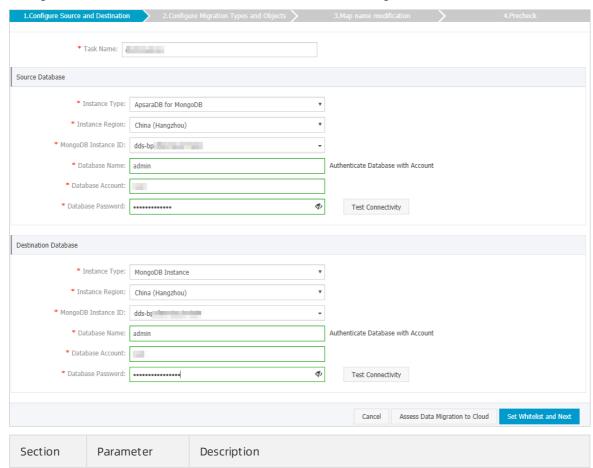
Create databases and collections to be sharded in the destination ApsaraDB for MongoDB instance, and configure data sharding based on your business requirements. For more information, see Configure sharding to maximize the performance of shards.

Note After you configure sharding for a cluster, data will not be migrated to the same shard.
This maximizes the performance of the sharded cluster.

- 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 for the data migration task.



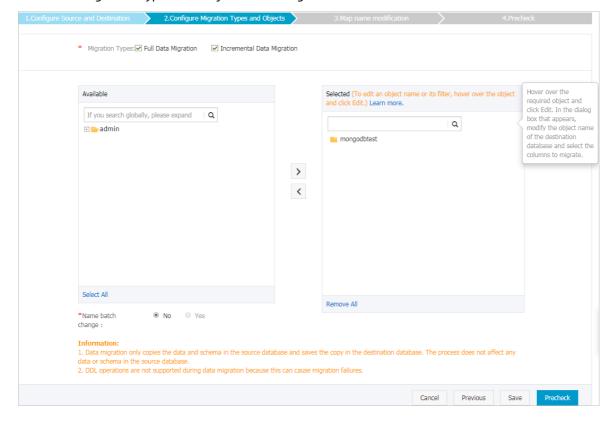
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select ApsaraDB for MongoDB.
	Inst ance Region	Select the region where the source ApsaraDB for MongoDB instance resides.
	MongoDB Instance ID	Select the ID of the source ApsaraDB for MongoDB instance.
		Enter the name of the authentication database. The database account is created in this database.
Source Dat abase	Dat abase Name	Note If the database account is root, enter admin.
	Dat abase Account	Enter the database account of the source ApsaraDB for MongoDB 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 source 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 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.
	Dat abase Name	Enter the name of the authentication database. The database account is created in this database.
		Note If the database account is root, enter admin.
Destinatio		

Pat abase Section	Parameter	Description
	Dat abase Account	Enter the database account of the destination ApsaraDB for MongoDB instance. For more information about the permissions that are required for the account, see Permissions required for database accounts.
		Enter the password of the destination 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, 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.

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



Paramet er	Description
Migratio	 To perform only full data migration, select only Full Data Migration. To ensure service continuity during data migration, select both Full Data Migration and Incremental Data Migration.
n Types	Note If Incremental Data Migration is not selected, do not write data to the source ApsaraDB for MongoDB database during full data migration. This ensures data consistency between the source and destination databases.
	 Select objects from the Available section and click the jecton icon to move the objects to the Selected section.
Objects	Note You cannot migrate data from the admin or local database.
	 You can select databases, collections, or functions as the objects to be migrated. After an object is migrated to the destination database, the name of the object remains unchanged. You can change the names of the objects that are migrated to the destination database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping.

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



325

- 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, click the icon next to each failed item to view details. Troubleshoot the issues based on the causes 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 and select Data Transmission Service (Pay-As-You-Go) Service Terms.
- 11. Click **Buy and Start** to start the migration task.
 - Full data migration
 We recommend that you do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the migration task automatically stops.
 - Increment al data migration
 An increment al data migration task does not automatically stop. You must manually stop the migration task.

- Note Select an appropriate time to manually stop the migration task. For example, you can stop the migration task during off-peak hours or before you switch your workloads to the destination ApsaraDB for MongoDB instance.
- 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. After the status of incremental data migration changes to The migration task is not delayed, manually stop the migration task.



12. Switch your workloads to the destination ApsaraDB for MongoDB instance.

9.9. Migrate the data of an ApsaraDB for MongoDB instance across regions

This topic describes how to migrate the data of a standalone instance or a replica set instance across regions by using Data Transmission Service (DTS). DTS supports both full data migration and incremental data migration. You can use these two methods together to migrate the data of an ApsaraDB for MongoDB instance across regions without interruptions to your business.

Prerequisites

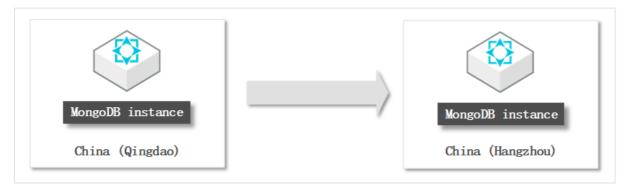
- The source instance is either a standalone instance or a replica set instance. If the source instance is a sharded cluster instance, we recommend that you use the built-in commands of MongoDB to migrate data. For more information, see Migrate a self-managed MongoDB database to ApsaraDB for MongoDB by using tools provided by MongoDB.
 - Note You cannot use DTS to incrementally migrate the data of 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 storage capacity of the destination instance must be greater than the occupied storage space of the source instance.

Context

You may need to migrate the data of an ApsaraDB for MongoDB instance across regions if:

- You want to restructure your business.
- You want to use the ApsaraDB for MongoDB instance to provide database services for applications deployed on an ECS instance, but the two instances are in different regions.

The following procedure illustrates how to migrate data from an ApsaraDB for MongoDB instance in China (Qingdao) to an instance in China (Hangzhou).



Note The procedure described in this topic only shows how to migrate the data of the source instance. If you no longer need the source instance after the migration is complete, you can release it.

Precautions

- We recommend that you migrate your data during off-peak hours to avoid business interruptions.
- To ensure data consistency, we recommend that you do not write data to the source instance while full data migration of a standalone instance is in progress.
- If the source and destination instances run different database versions or storage engines, ensure there are no compatibility issues between them before you start migration. For more information about the database versions and storage engines 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	Note 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	 Note The create and delete operations that are performed on databases, collections, and indexes can be synchronized. The create, delete, and update operations that are performed on documents can be synchronized. 		

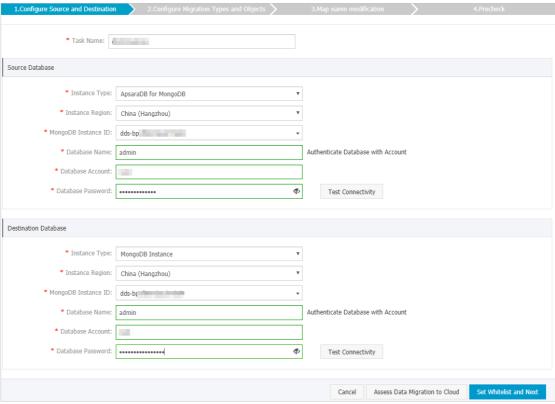
Required database account permissions

Data source	Full data migration	Incremental data migration
Source ApsaraDB for MongoDB instance	Read permissions on the source database	Read permissions on the source database, admin database, and local database
Destination ApsaraDB for MongoDB instance	Read/write permissions on the destination database	Read/write permissions on the destination database

Note For more information about how to create and authorize a database account, see Use DMS to manage MongoDB users.

Procedure

- 1.
- 2.
- 3. In the upper-right corner of the page that appears, click Migrate Databases.
- 4. On the Create Migration Task page, create a migration task.
 - i. Configure the source and destination databases.



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	Note If you want to use the root account, enter admin in the Database Name field.	
		The username of the account that is used to connect to the source	
Source	Dat abase Account	instance. For information about the permissions that are required for the account, see .	
Database			

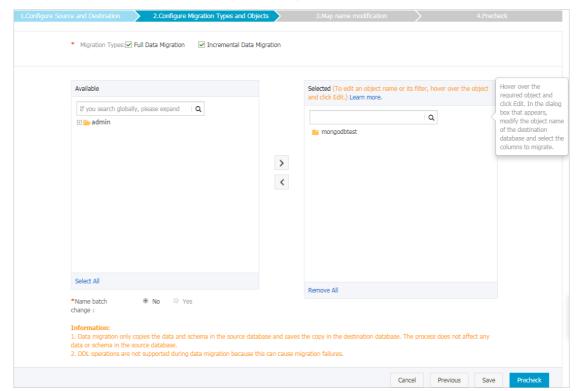
Section	Parameter	Parameter		
		The password of the database account.		
	Database Password	? Note After you specify the information about the source database, you can click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed and you must click Check next to the Failed message to modify the information.		
	Instance Type	Select MongoDB Instance.		
	Instance Region	The region where the destination ApsaraDB for MongoDB instance 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.		
Destinati		Note If you want to use the root account, enter admin in the Database Name field.		
on Database	Database 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 .		
	Dat abase Password	The password of the database account.		
		Note After you specify the information about the destination database, you can click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message appears. If the information is incorrect, the Failed message appears and you must click Check next to the Failed message to modify the information.		

ii. 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.





Setting	Description
Select migrati on types	 If you want to migrate only the existing data in the source instance, select Full Data Migration. If you want to migrate data without downtime, select Full Data Migration and Incremental Data Migration. 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.
	The method to select objects to migrate: a. Select one or more objects from the Available section. b. Click > to add the selected objects to the Selected section.
Select objects to migrate	 Note DTS cannot migrate data from the admin or local database. The objects that you can migrate are databases, collections, and functions. 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.
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	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. 7 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 prode. You can also release the DTS instance at your partiest expectable after the
or destinat ion databas e	needs. You can also release the DTS instance at your earliest opportunity after the source and destination instances are released.

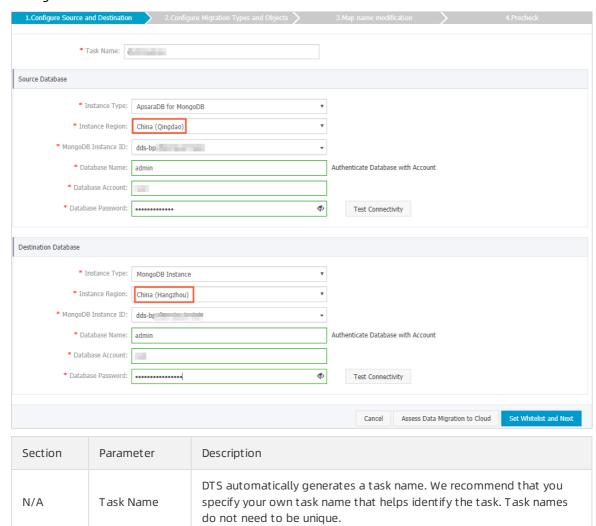
iv. In the lower-right corner of the page, click Precheck.



- 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.
- 5. Configure both the source and destination databases.

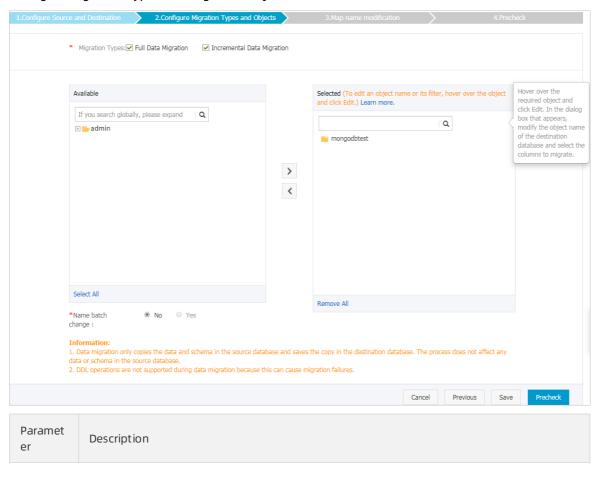


Section	Parameter	Description		
	Instance Type	Select ApsaraDB for MongoDB.		
	Instance Region	Select the region where the source ApsaraDB for MongoDB instance resides. For this example, select China (Qingdao) .		
	MongoDB Instance ID	Select the ID of the source ApsaraDB for MongoDB instance.		
	Database	Enter the name of the authentication database. It is the database where the database account is created.		
Source	Name	Note If the database account is root, enter admin.		
Database	Dat abase Account	Enter the username of the database account you use to manage the source database. For more information about the account permission requirements, see Required database account permissions.		
		Enter the password of the database account.		
	Dat abase Password	Note After you specify the source database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed, and you must click Check next to the Failed message to modify the information as prompted.		
	Instance Type	Select MongoDB Instance.		
	Instance Region	Select the region where the destination ApsaraDB for MongoDB instance resides. For this example, select China (Hangzhou) .		
	MongoDB Instance ID	Select the ID of the destination ApsaraDB for MongoDB instance.		
	Database Name	Enter the name of the authentication database. It is the database where the database account is created.		
		? Note If the database account is root, enter admin.		
Destinatio	Dat abase Account	Enter the username of the database account you use to manage the destination database. For more information about the account permission requirements, see Required database account permissions.		
n Dat abase				

Section Par	arameter	Description	
		Enter the password of the database account.	
	atabase assword	? Note After you specify the destination database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed, and you must click Check next to the Failed message to modify the information as prompted.	

6.

7. Configure migration types and migration objects.



Paramet er	Description
	 If you want to migrate all data, select Full Data Migration. If you want to migrate data without interruptions to your business, select both Full Data Migration and Incremental Data Migration.
Migratio n Types	 Note You cannot use DTS to incrementally migrate the data of a standalone instance. If Incremental Data Migration is not selected, do not write data into the source database during full data migration. This ensures data consistency between the source and destination databases.
	 In the Available section, select the objects you want to migrate and then click the joint of move them to the Selected section.
Availabl e	Note Data in the admin database cannot be migrated even if this database is selected.
	 A migration object can be a database, collection, or function. By default, the name of an object remains unchanged after migration. If you want a different object name after migration, use the object name mapping feature provided by DTS. For more information, see Object name mapping.

8.

9.

10.

- 11. View the progress of the migration task.
 - Full dat a 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.

 Increment al data migration
 An increment al 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.

What to do next

Determine whether to release the source instance.

- If the source instance uses pay-as-you-go billing, release it. For more information, see Release an instance or a node.
- If the source instance uses subscription billing, you cannot release it.

10.Migrate across different Alibaba Cloud accounts

10.1. Migrate data between Apsara PolarDB for MySQL clusters across Alibaba Cloud accounts

Apsara PolarDB is a next-generation relational database service developed by Alibaba Cloud. It is a high-performance, high-availability, easy-to-use, and reliable service that is compatible with the MySQL database engine. This topic describes how to migrate data between Apsara PolarDB for MySQL clusters that are owned by different Alibaba Cloud accounts by using Data Transmission Service (DTS).

Prerequisites

- The source and destination Apsara PolarDB for MySQL clusters are created. For more information, see Create a PolarDB MySQL cluster.
- The binary logging feature is enabled for the source Apsara 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.
- 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 Alibaba Cloud over the Internet. For more
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	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE 	

Permissions required for database accounts

Database	Required permission
Source Apsara PolarDB for MySQL cluster	The read permission on the objects to be migrated
Destination Apsara PolarDB for MySQL cluster	The read/write permissions on the objects to be migrated

Note For more information about how to create and authorize a database account, see Create database accounts.

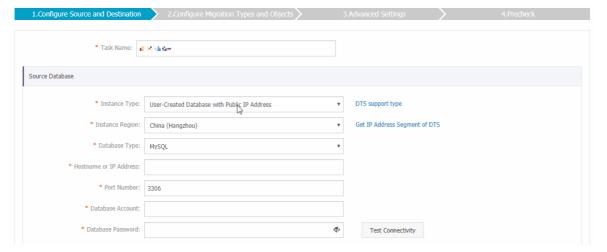
Preparations

- 1. Log on to the Alibaba Cloud console by using the Alibaba Cloud account to which the source Apsara PolarDB for MySQL cluster belongs.
- 2. Create a RAM role and authorize this role to access the cloud resources of the Alibaba Cloud account to which the source Apsara PolarDB for MySQL cluster belongs. 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.

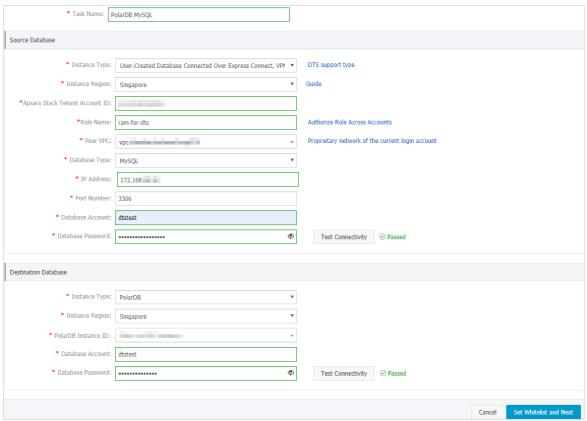
Note To migrate data between Apsara PolarDB for MySQL clusters across Alibaba Cloud accounts, you can specify the source Apsara PolarDB for MySQL cluster as a user-created database connected over Express Connect. You can also request a public endpoint for the source Apsara PolarDB for MySQL cluster and specify the source cluster as a user-created database with a public IP address. In this case, you do not need to create a RAM role and configure authorization.

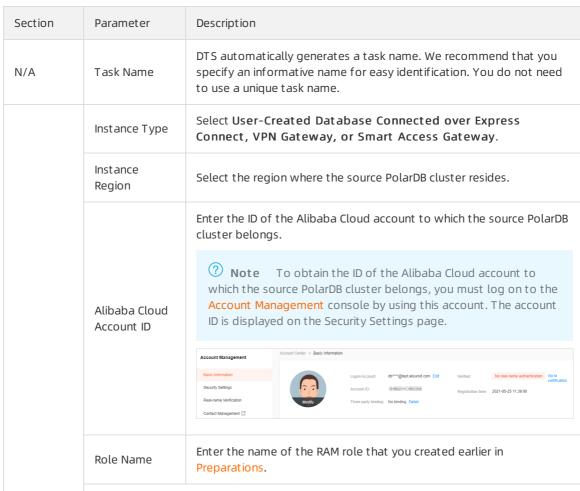
Procedure

- 1. Log on to the DTS console by using the Alibaba Cloud account to which the destination Apsara PolarDB for MySQL cluster belongs.
- 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 Proprietary network of Other Apsara Stack Accounts next to the Peer VPC field.
 - Note To migrate data between Apsara PolarDB for MySQL clusters across Alibaba Cloud accounts, you can specify the source Apsara PolarDB for MySQL cluster as a user-created database connected over Express Connect. You can also request a public endpoint for the source Apsara PolarDB for MySQL cluster and specify the source cluster as a user-created database with a public IP address.



6. Configure the source and destination databases for the data migration task.





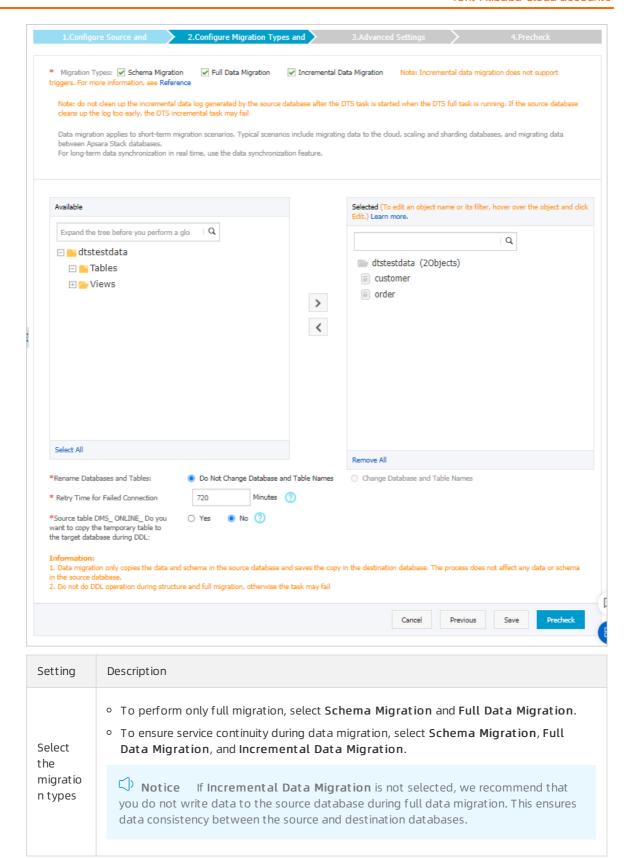
Peer VPC	Select the ID of the VPC where the source PolarDB cluster resides. To obtain the VPC ID, you must log on to the Apsara PolarDB console by using the Alibaba Cloud account to which the source PolarDB cluster belongs. On the Clusters page, click the ID of the source PolarDB cluster. The VPC ID is displayed in the Basic Information section. (Cluster	
	Maintenance Window 02:00-03:00 Modify	
Database Type	Select MySQL.	
IP Address	Enter the private IP address of the source PolarDB cluster. In this example, enter 172.16.20.20. You can obtain the private IP address by pinging the VPC-facing endpoint of the source PolarDB cluster. Connection Information © Vertical Configure Delate Settings and Manag. Diagnostics and Opt. VPC-tacrog Endpoint VPC-tacrog Endpoint VPC-tacrog Endpoint VPC-tacrog Endpoint Public-facing Endpoint VPC-tacrog Endpoint V	
Port Number	Enter the service port number of the source PolarDB cluster. The default port number is 3306 .	
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.	
Dat abase Password	Procession of the source database account. 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.	
	Port Number Database Account	

Section	Parameter	Description	
Destinatio n Database	Instance Type	Select PolarDB.	
	Instance Region	Select the region where the destination PolarDB cluster resides.	
	PolarDB Instance ID	The ID of the destination PolarDB cluster.	
	Dat abase Account	Enter 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.	
	Dat abase Password	Enter the password of the destination database account.	
		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.	

7. In the lower-right corner of the page, click \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{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.



Setting	Description		
Select the objects to be migrate d	Select one or more objects from the Available section and click the icon to move the objects to the Selected section. Notice You can select columns, tables, or databases as the objects to be migrated. 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. 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 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.		
	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.		

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	Note If online DDL operations generate a large amount of data, the data migration task may be delayed.
perform s online DDL operatio ns on the source table	• No : 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.
	Note If you select No, the tables in the destination database may be locked.

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

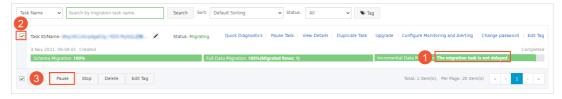


- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.
- \circ If the task fails to pass the precheck, you can click the \bigcirc icon next to each failed item

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.
- 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/write permissions. After the data migration is complete, you must delete the database accounts to ensure security.

10.2. Migrate data between ApsaraDB for MongoDB instances created by different Alibaba Cloud accounts

This topic describes how to migrate data between ApsaraDB for MongoDB instances created by different Alibaba Cloud accounts by using Data Transmission Service (DTS). DTS supports both full data migration and incremental data migration. You can use these two methods together to migrate data between ApsaraDB for MongoDB instances without interruptions to your business.

Prerequisites

- The source instance is either 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 commands 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 You cannot use DTS to incrementally migrate the data of 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 storage capacity of the destination instance must be greater than the occupied storage space of 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		
	DTS migrates the historical data of required objects from the source MongoDB database to the destination MongoDB database.		
Full data migration	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	 Note The create and delete operations that are performed on databases, collections, and indexes can be synchronized. The create, delete, and update operations that are performed on documents can be synchronized. 		

Required database account permissions

Data source	Full data migration	Incremental data migration
Source ApsaraDB for MongoDB instance	Read permissions on the source database	Read permissions on the source database, admin database, and local database
Destination ApsaraDB for MongoDB instance	Read/write permissions on the destination database	Read/write permissions on the destination database

Note For more information about how to create and authorize a database account, see Use DMS to manage MongoDB users.

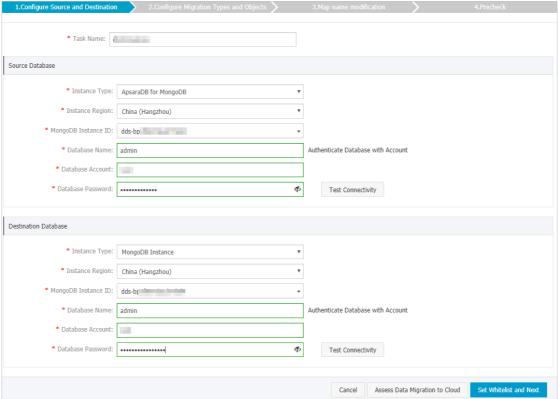
Preparation

- 1. Log on to the ApsaraDB for MongoDB console with the Alibaba Cloud account to which the source instance belongs.
- 2. Apply for a public endpoint for the source instance. For more information, see Apply for a public endpoint.
- 3. Add the Classless Inter-Domain Routing (CIDR) blocks of DTS servers to a 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 IP whitelist of on-premises databases. For example, if the source instance is in China (Hangzhou) and the destination instance is in China (Shenzhen), add the CIDR blocks of the DTS servers in China (Shenzhen) to a whitelist of the source instance.

Procedure

- 1. Log on to the DTS console with the Alibaba Cloud account to which the destination ApsaraDB for MongoDB instance belongs.
- 2.
- 3. In the upper-right corner of the page that appears, click Migrate Databases.
- 4. On the Create Migration Task page, create a migration task.
 - i. Configure the source and destination databases.



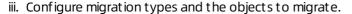
Section	Parameter	Parameter
N/A	Task Name	The task name that DTS automatically generates. We recommend that you specify a descriptive name that makes it easy to identify the task. You do not need to specify a unique task name.
	Instance Type	Select ApsaraDB for MongoDB.
	Instance Region	The region where 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.
Database Name		Note If you want to use the root account, enter admin in the Database Name field.
	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 .
Source Database		

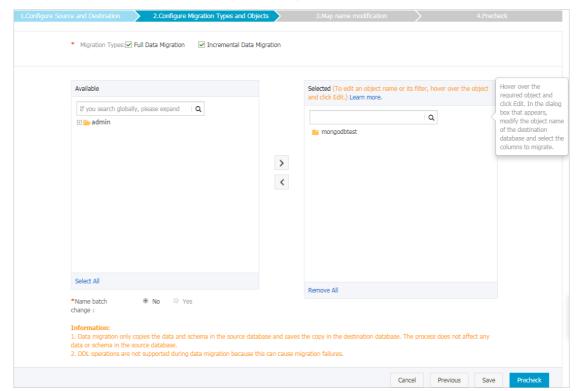
Section	Parameter	Parameter	
	Dat abase Password	Note After you specify the information about the source database, you can click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed and you must click Check next to the Failed message to modify the information.	
	Instance Type	Select MongoDB Instance.	
	Instance Region	The region where the destination ApsaraDB for MongoDB instance resides.	
Destinati on Database	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. Note If you want to use the root account, enter admin in the Database Name field.	
	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 .	
	Dat abase Password	The password of the database account. Note After you specify the information about the destination database, you can click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message appears. If the information is incorrect, the Failed message appears and you must click Check next to the Failed message to modify the information.	

ii. 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.





Setting	Description
Select migrati on types	 If you want to migrate only the existing data in the source instance, select Full Data Migration. If you want to migrate data without downtime, select Full Data Migration and Incremental Data Migration. 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.
Solort	The method to select objects to migrate: a. Select one or more objects from the Available section. b. Click to add the selected objects to the Selected section. 7 Note
Select objects to migrate	 DTS cannot migrate data from the admin or local database. The objects that you can migrate are databases, collections, and functions. 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.
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 ion	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. ② 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.
databas e	

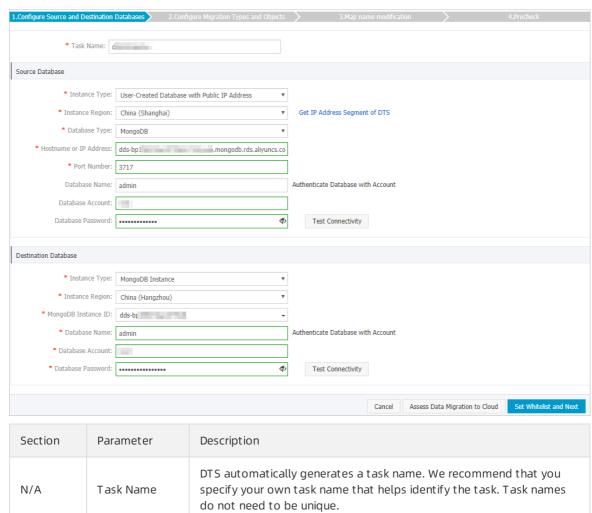
iv. In the lower-right corner of the page, click Precheck.



- 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.
- 5. Configure both the source and destination databases.



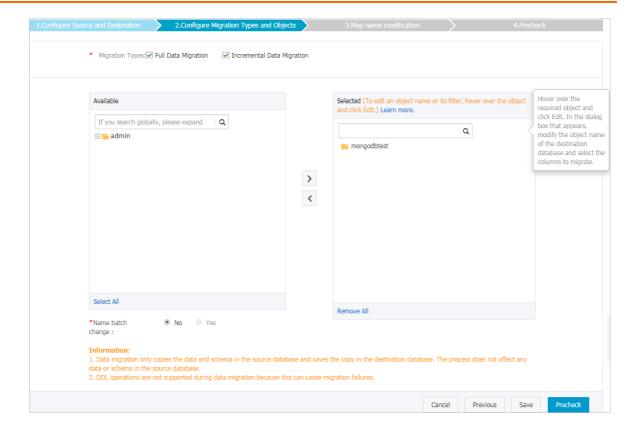
		Description
Source Database	Instance Type	Select User-Created Database with Public IP Address.
	Instance Region	If you set the instance type to User-Created Database with Public IP Address, the system automatically specifies Instance Region .
	Database Type	Select MongoDB.
	Hostname or IP Address	Enter 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 3717, which is the service port of the source instance.
	Database Name	Enter the name of the authentication database. It is the database where the database account is created.
		Note If the database account is root, enter admin.
	Dat abase Account	Enter the username of the database account you use to manage the source database. For more information about the account permission requirements, see Required database account permissions.
	Dat abase Password	Note After you specify the source database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed, and you must click Check next to the Failed message to modify the information as prompted.
	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.

Section	Parameter	Description
Destinatio n Database	Database Name	Enter the name of the authentication database. It is the database where the database account is created.
		? Note If the database account is root, enter admin.
	Database Account	Enter the username of the database account you use to manage the destination database. For more information about the account permission requirements, see Required database account permissions.
	Dat abase Password	Enter the password of the database account.
		Note After you specify the destination database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed, and you must click Check next to the Failed message to modify the information as prompted.
	Account	destination database. For more information about the account permission requirements, see Required database account permission requirements, see Required database account permission. The Password of the database account. Note After you specify the destination database information, click Test Connectivity next to Database Password to check whether the information is correct. If the information is correct, the Passed message is displayed. If the information is incorrect, the Failed message is displayed, and must click Check next to the Failed message to modify the

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

Note The IP addresses of DTS servers are automatically added to a whitelist of the destination instance. This ensures that the DTS servers can connect to the destination instance. After the migration is complete, you can remove the IP addresses from the whitelist if you no longer need them. For more information, see Configure a whitelist.

7. Configure migration types and migration objects.



- 8.
- 9.
- 10.
- 11. View the progress of the migration task.
 - Full dat a 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.

References

If you migrate data to a sharded cluster instance, you can configure data sharding as needed. For more information, see Configure sharding to maximize the performance of shards.

11.Migrate from Alibaba Cloud to a user-created database

11.1. Migrate data from an ApsaraDB RDS for MySQL database to a user-created MySQL database

This topic describes how to migrate data from an ApsaraDB RDS for MySQL database to a user-created 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 database contain primary keys or UNIOUE NOT NULL indexes.
- The available storage space of the user-created MySQL database is greater than the total size of the data in the ApsaraDB RDS for MySQL database.
- The version of the user-created MySQL database is the same as that of the ApsaraDB RDS for MySQL database. This ensures compatibility.

SQL operations that can be synchronized during incremental data migration

Permissions required for database accounts

Database	Permission
ApsaraDB RDS for MySQL database	The read permission for the objects to be migrated
User-created MySQL database	The read/write permissions for the objects to be migrated

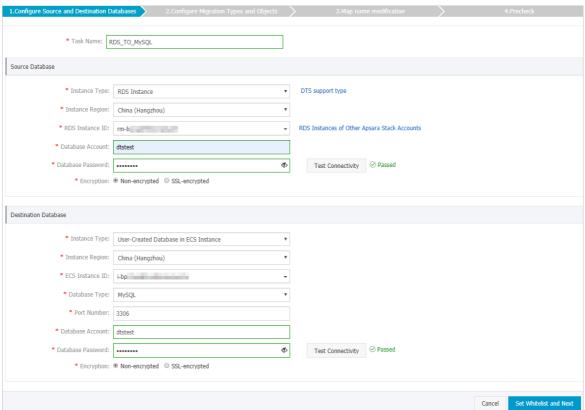
For more information about how to create and authorize a database account, see the following topics:

- Create an account for an RDS for MySQL instance and Change the permissions of an account for an RDS for MySQL instance
- Create an account for a user-created MySQL database and configure binary logging for a user-created MySQL database

Procedure

- 1.
- 2.
- 3.
- 4.

5. Configure the information about the source and destination databases for the data migration task.



Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you use an informative name for easy identification. You do not need to use a unique task name.
	Instance Type	Select RDS Instance.
	Instance Region	Select the region where the source RDS instance resides.
	RDS Instance ID	Select the ID of the source RDS instance.
	Dat abase Account	Enter the database account of the source RDS instance. For more information about permissions required for the account, see Permissions required for database accounts.

Caurca

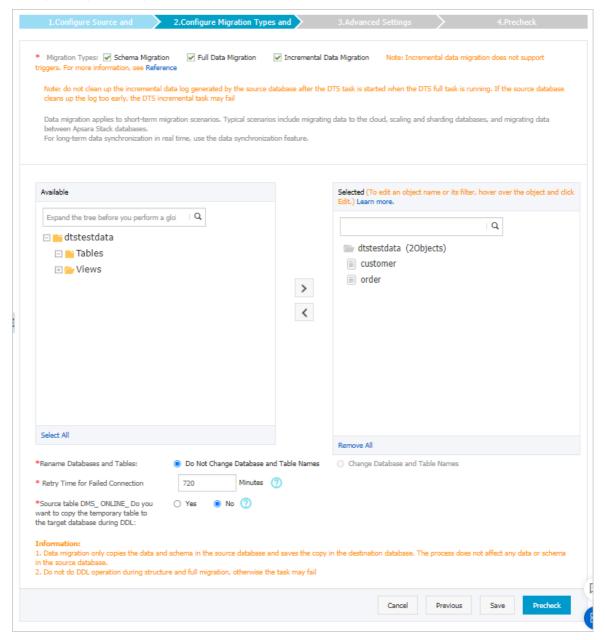
Source Dettiba se	Parameter	Description
	Database Password	Pote After the source database information is specified, click Test Connectivity next to Database Password to verify whether the specified information is valid. If the specified information is valid, the Passed message appears. If the Failed message appears, click Check in the Failed message. Modify the source database information based on the instructions.
	Encryption	Select Non-encrypted or SSL-encrypted. If you want to select SSL-encrypted, you must enable SSL encryption for the RDS instance before configuring the data migration task. For more information, see Configure SSL encryption for an RDS for MySQL instance. Note The Encryption parameter is available only in mainland China and Hong Kong(China).
	Instance Type	Select User-Created Database in ECS Instance.
	Instance Region	Select the region where the ECS instance resides.
	ECS Instance	Select the ID of the ECS instance that is connected to the user-created MySQL database.
	Dat abase Type	Select MySQL.
	Port Number	Enter the service port number of the user-created MySQL database. In this example, enter 3306 .
Dockinskins	Dat abase Account	Enter the account of the user-created MySQL database. For more information about permissions required for the account, see Permissions required for database accounts.
Destination Database	Dat abase Password	Enter the password for the account of the user-created MySQL database. ? Note After the destination database information is specified, click Test Connectivity next to Database Password to verify whether the specified information is valid. If the specified information is valid, the Passed message appears. If the Failed message appears, click Check in the Failed message. Modify the destination database information based on the instructions.

Section	Parameter	Description
	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 The CIDR blocks of DTS servers are automatically added to the whitelist of the source RDS instance and the inbound rule of the destination ECS instance. This ensures that DTS servers can connect to the source and destination instances.

7. Configure migration types and objects.



ltem	Description			
Migratio n types	 To perform only full data migration, select Schema Migration and Full Data Migration. If you want to migrate data without disruptions to your business, select Schema Migration, Full Data Migration, and Incremental Data Migration. Note If Incremental Data Migration is not selected, do not write data into the source database during full data migration. This ensures data consistency between the source and destination databases. 			
	Select the objects to be migrated in the Available section and click icon to move them to the Selected section.			
Objects to be migrate d	 Note Objects to be migrated can be databases, tables, or columns. The selected objects are not renamed after the migration by default. If you want to rename the objects migrated to the destination instance, you can use the object name mapping feature provided by DTS. For more information, see Object name mapping. If you use the object name mapping feature for an object, objects that depend on the object may fail to be migrated. 			

8.

9.

10.

11.

11.2. Migrate data from an Apsara PolarDB for MySQL cluster to a user-created MySQL database

This topic describes how to migrate data from an Apsara PolarDB for MySQL cluster to a user-created 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 for the Apsara PolarDB for MySQL cluster is enabled. For more information, see Enable binary logging.

Precautions

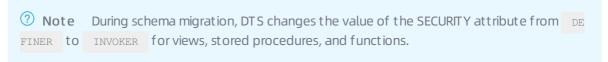
• DTS uses read and write resources of the source and destination databases during full data migration. This may increase the loads of the database servers. If the database performance is

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.
- 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.



• 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 Alibaba Cloud over the Internet. For more	
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	 ALTER TABLE and ALTER VIEW CREATE FUNCTION, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, and CREATE VIEW DROP INDEX and DROP TABLE RENAME TABLE TRUNCATE TABLE

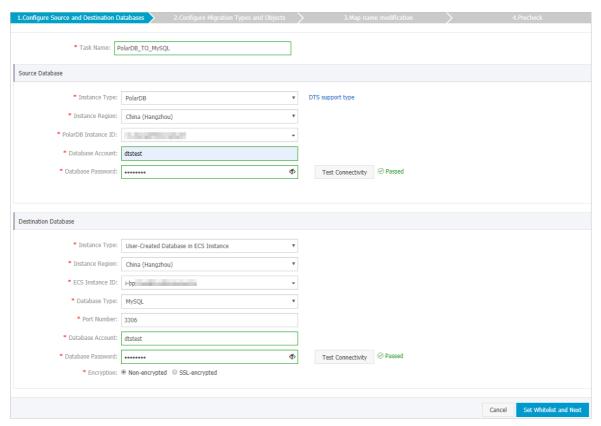
Permissions required for database accounts

Database	Permission	
Apsara PolarDB for MySQL cluster	The read permission for the objects to be migrated	
User-created MySQL database	The read/write permissions for 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 an RDS for MySQL instance.

Procedure

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Dat a 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.



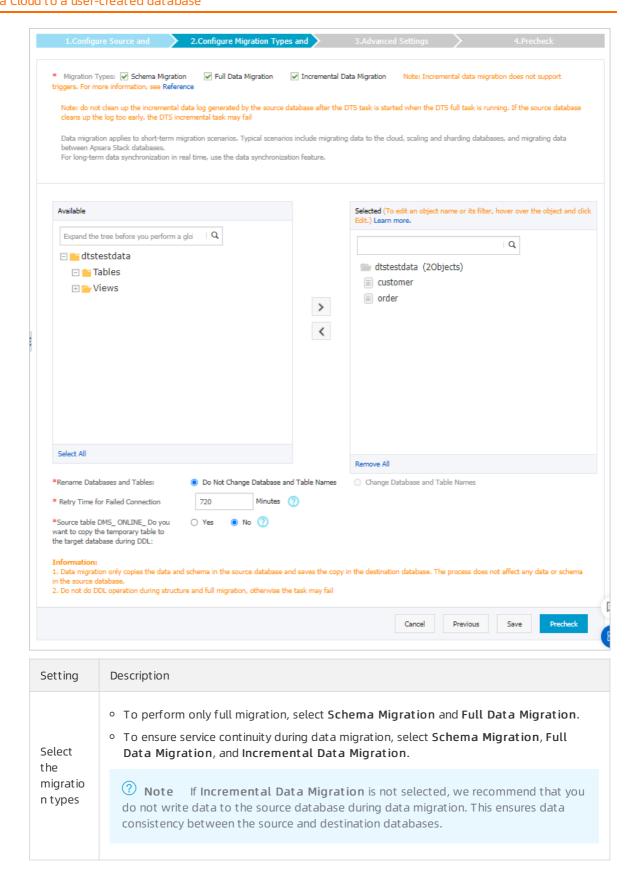
Section	Parameter	Description		
N/A	Task Name	DTS automatically generates a task name. We recommend that you use an informative name for easy identification. You do not need to use a unique task name.		
	Instance Type	Select PolarDB .		
	Instance Region	Select the region where the source PolarDB cluster resides.		
	PolarDB Instance ID	Select the ID of the source PolarDB cluster.		
	Dat abase Account	Enter the database account for the source PolarDB cluster. For more information about permissions required for the account, see Permissions required for database accounts.		
Source Database		Enter the password for the database account.		
	Dat abase Password	Note After you specify the source database parameters, click Test Connectivity next to the Database Password parameter 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 User-Created Database in ECS Instance.	
	Instance Region	Select the region where the ECS instance resides.	
	ECS Instance	Select the ID of the ECS instance that is connected to the user-created MySQL database.	
	Dat abase Type	Select MySQL.	
	Port Number	Enter the service port number of the user-created MySQL database. In this example, enter 3306 .	
Destination Database	Dat abase Account	Enter the account for the user-created MySQL database. For more information about permissions required for the database account, see Permissions required for database accounts.	
		Enter the password for the database account.	
	Dat abase Password	Note After you specify the destination database parameters, click Test Connectivity next to the Database Password parameter 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.	
	Encryption	Select Non-encrypted or SSL-encrypted. In this example, Non-encrypted is selected.	

6. Click Set Whitelist and Next in the lower-right corner of the page.

? Note The CIDR blocks of DTS servers are automatically added 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 destination instance.

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



Setting	Description
Select the objects to be migrate d	Select one or more objects from the Available section and click the icon to move the objects to the Selected section. Note You can select columns, tables, or databases as the objects to be migrated. 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. 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 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 e	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.

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	Note If online DDL operations generate a large amount of data, the data migration task may be delayed.		
	 No: 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. 		
	? Note If you select No, the tables in the destination database may be locked.		
source table			

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



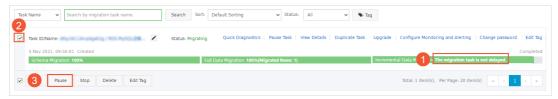
369

- 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.
- \circ If the task fails to pass the precheck, you can click the or 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.
- 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.Migrate between usercreated databases

12.1. Migrate data between usercreated Oracle databases

This topic describes how to migrate data between user-created 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 the versions of the source and destination databases are the same.

- Supplemental logging, including SUPPLEMENTAL_LOG_DATA_PK and SUPPLEMENTAL_LOG_DATA_UI, is enabled for the source Oracle database. For more information, see Supplemental Logging.
- The ARCHIVELOG mode is enabled for the source Oracle database. 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.
- DTS automatically resumes a failed data migration 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 from Alibaba Cloud over the Internet. For more
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.		
3	Note If an object contains triggers, the data between the source and destination databases will become inconsistent.		
	DTS migrates historical data of the required objects from the source Oracle database to the destination Oracle database.		
Full data migration	Note 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.		
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 owner permission on schemas	The owner permission on schemas	The database administrator (DBA) permission
Destination Oracle database	The owner permission on schemas	The owner permission on schemas	The owner permission on schemas

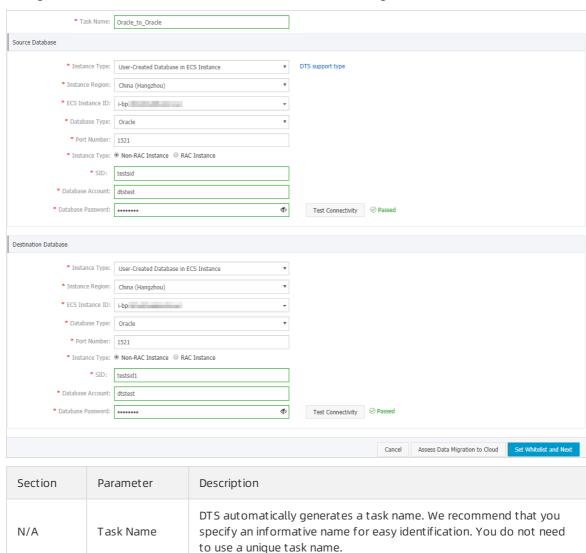
For more information about how to create and authorize an Oracle database account, see CREATE USER and GRANT.

Notice If you want 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

The procedure in this topic uses a **user-created database hosted on ECS** as an example. You can also follow the procedure to configure data migration tasks for other types of user-created Oracle databases.

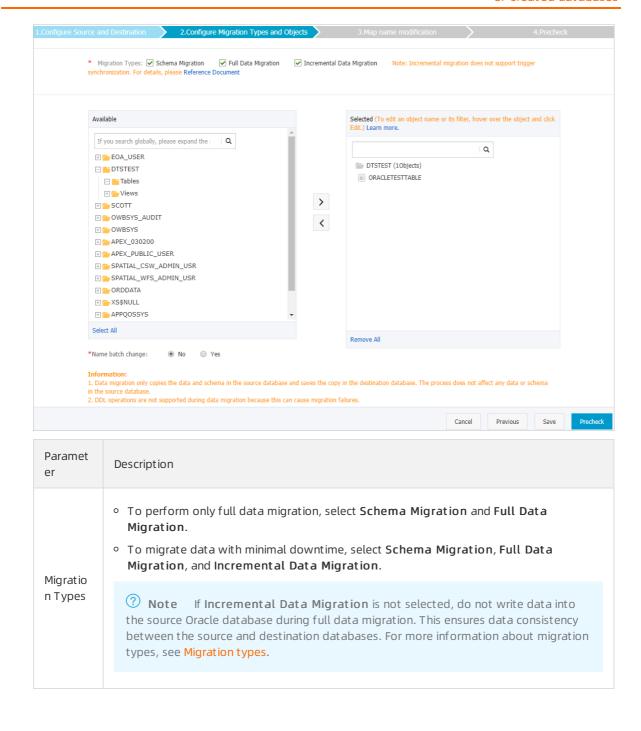
- 1.
- 2.
- 3.
- 4.
- 5. Configure the source and destination databases for the data migration task.



Section	Parameter	Description
Source Database	Instance Type	Select User-Created Database in ECS Instance. Note If you select other instance types, you must prepare the environments that are required for the source database. For more information, see Preparation overview.
	Inst ance Region	Select the region of the 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.
	Instance Type	 Non-RAC Instance: If you select this option, you must specify the SID. RAC Instance: If you select this option, you must specify the Service Name.
	Dat abase Account	Enter the account of the source Oracle database. For more information about the permissions that are required for the account, see Permissions required for database accounts.
	Database Password	Enter the password for the source database account.
		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 User-Created Database in ECS Instance.
	Inst ance 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.
	Dat abase Type	Select Oracle .
		Enter the service port number of the destination Oracle database.

Section	Parameter	Description
Destinatio n Database	Instance Type	 Non-RAC Instance: If you select this option, you must specify the SID. RAC Instance: If you select this option, you must specify the Service Name.
	Dat abase Account	Enter the account of the destination Oracle 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 for the destination database account.
		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 \mathbf{Set} $\mathbf{Whitelist}$ and \mathbf{Next} .
- 7. Select the migration types and objects to be migrated.



Paramet er	Description		
	Select objects from the Available section and click the icon to move the objects to the Selected section.		
Objects	 Note You can select columns, tables, or databases as the objects to be migrated. 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 change the names of the objects that are migrated to the destination Oracle database by using the object name mapping feature. For more information about how to use this feature, see Object name mapping. If you use the object name mapping feature on an object, other objects that are dependent on the object may fail to be migrated. 		

8.

9.

10.

11.