

Alibaba Cloud Data Transmission Service

Data migration

Issue: 20200610

Legal disclaimer

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

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

- 5.** By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- 6.** Please contact Alibaba Cloud directly if you discover any errors in this document.

Document conventions

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

Style	Description	Example
{ } or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

Contents

Legal disclaimer	1
Document conventions	1
1 Overview of data migration scenarios	1
2 Database and migration types	5
3 Data type mappings between heterogeneous databases	15
4 Migration task management	21
4.1 Object name mapping.....	21
4.2 Scheduling policies of regular data migration.....	29
4.3 Use conditions to filter data.....	31
4.4 View the progress of a data migration task.....	35
4.5 View the performance of full data migration.....	38
4.6 Check the performance of incremental data migration.....	39
4.7 Adjust the full migration rate.....	40
4.8 Fix a failed data migration task.....	42
4.10 Stop a data migration task.....	44
5 Precheck and troubleshooting	47
5.1 Source database connectivity.....	47
5.2 Source database version.....	49
6 Migration from a user-created database to Alibaba Cloud	50
6.1 Migrate data from a user-created MySQL database to an ApsaraDB RDS for MySQL instance.....	50
6.2 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.....	59
6.3 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.....	69
6.4 Migrate data from a user-created MySQL database to an ApsaraDB for PolarDB cluster.....	80
6.5 Migrate data from a user-created MySQL database to a DRDS instance.....	88
6.6 Migrate incremental data from a user-created SQL Server database to an ApsaraDB RDS for SQL Server database.....	97
6.7 Migrate full data from a user-created SQL Server database to ApsaraDB RDS for SQL Server.....	109
6.8 Migrate data from a user-created Oracle database to a PolarDB cluster compatible with Oracle.....	117
6.9 Migrate data from a user-created Oracle database to an ApsaraDB RDS for MySQL instance.....	125
6.10 Migrate data from a user-created Oracle database to a DRDS instance.....	136

6.11 Migrate data from a user-created Oracle database to an ApsaraDB RDS for PPAS instance.....	145
6.12 Migrate incremental data from a user-created PostgreSQL database (version 10.x to 12) to an ApsaraDB RDS for PostgreSQL instance.....	155
6.13 Migrate incremental data from a user-created PostgreSQL database (version 9.4 to 9.6) to an ApsaraDB RDS for PostgreSQL instance.....	164
6.14 Migrate full data from a user-created PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance.....	176
6.15 Migrate data from a user-created Redis database to an ApsaraDB for Redis instance.....	184
6.16 Migrate user-created standalone MongoDB databases to Alibaba Cloud by using DTS.....	193
6.17 Migrate the replica set of a user-created MongoDB database to ApsaraDB for MongoDB by using DTS.....	201
6.18 Migrate the user-created sharded MongoDB database to ApsaraDB for MongoDB by using DTS.....	210
6.19 Migrate incremental data from a user-created TiDB database to an ApsaraDB RDS for MySQL instance.....	220
6.20 Migrate full data from a user-created TiDB database to an ApsaraDB RDS for MySQL instance.....	239
6.21 Migrate data from a user-created Db2 database to an ApsaraDB RDS for MySQL database.....	248
7 Migration from a third-party cloud to Alibaba Cloud.....	257
7.1 Migrate data from an Amazon RDS for MySQL database to an ApsaraDB RDS for MySQL database.....	257
7.2 Migrate data from an Amazon RDS for Oracle database to an ApsaraDB RDS for MySQL database.....	267
7.3 Migrate full data from an Amazon RDS for Oracle database to an ApsaraDB RDS for PPAS database.....	278
7.4 Migrate incremental data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL database.....	287
7.5 Migrate full data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL database.....	297
7.6 Migrate data from an Amazon Aurora MySQL database to an ApsaraDB RDS for MySQL database.....	306
7.7 Migrate data from an Amazon Aurora MySQL database to a PolarDB for MySQL database.....	317
7.8 Migrate full data from an Amazon Aurora PostgreSQL database to an ApsaraDB RDS for PostgreSQL database.....	328
7.9 Migrate full data from an Amazon RDS for SQL Server database to an ApsaraDB RDS for SQL Server database.....	338
7.10 Migrate a MongoDB Atlas database to Alibaba Cloud through DTS.....	347
8 Migrate data between instances under the same Alibaba Cloud account.....	356

8.1 Migrate data between RDS instances.....	356
8.2 Migrate full data between ApsaraDB RDS for MySQL instances on a regular basis.....	364
8.3 Migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance.....	373
8.4 Migrate data from an ApsaraDB RDS for MySQL database to an ApsaraDB for PolarDB database.....	382
8.5 Migrate data between ApsaraDB for PolarDB clusters.....	391
8.6 Migrate data from an ApsaraDB for PolarDB database to an ApsaraDB RDS for MySQL database.....	398
8.7 Migrate data between PolarDB clusters compatible with Oracle.....	407
8.8 Migrate data from a standalone instance to a replica set instance or a sharded cluster instance.....	415
8.9 Migrate data from a replica set instance to a sharded cluster instance.....	420
8.10 Migrate the data of an ApsaraDB for MongoDB instance across regions.....	427
9 Migrate data across Alibaba Cloud accounts.....	434
9.1 Migrate data between Apsara PolarDB for MySQL clusters across Alibaba Cloud accounts.....	434
9.2 Migrate data between ApsaraDB for MongoDB instances created by different Alibaba Cloud accounts.....	444
9.3 Migrate data between ApsaraDB for RDS instances created by different Alibaba Cloud accounts.....	450
10 Migration from Alibaba Cloud to a user-created database... 457	
10.1 Migrate data from an ApsaraDB RDS for MySQL database to a user-created MySQL database.....	457
10.2 Migrate data from an Apsara PolarDB for MySQL cluster to a user-created MySQL database.....	465
11 Migrate data between user-created databases.....	474
11.1 Migrate data between user-created Oracle databases.....	474

1 Overview of data migration scenarios

You can use DTS to migrate data between homogeneous and heterogeneous data sources. DTS provides the following extract, transform, and load (ETL) capabilities: object name mapping for columns, tables, and databases, and data filtering. You can migrate data to or from Alibaba Cloud. You can also migrate data between instances that are provided by Alibaba Cloud.



Note:

For more information about the databases supported by DTS, see [Database and migration types](#).

Scenario	Source database	Topic
Migrate data from a user-created database to Alibaba Cloud	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 database
		Migrate data from a user-created MySQL database to an ApsaraDB for PolarDB cluster
		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 database
		Migrate full data from a user-created SQL Server database to ApsaraDB RDS for SQL Server

Scenario	Source database	Topic
	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
	PostgreSQL	Migrate incremental data from a user-created PostgreSQL database (version 10.x to 12) to an ApsaraDB RDS for PostgreSQL instance
		Migrate incremental data from a user-created PostgreSQL database (version 9.4 to 9.6) 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 an ApsaraDB for Redis instance
	MongoDB	Migrate data from a standalone user-created MongoDB instance to Alibaba Cloud
		Migrate data from a user-created MongoDB instance in the replica set architecture to Alibaba Cloud
		Migrate data from a user-created MongoDB instance in the sharded cluster architecture to Alibaba Cloud
	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

Scenario	Source database	Topic	
	Db2	Migrate data from a user-created Db2 database to an ApsaraDB RDS for MySQL database	
Migrate data from a third-party cloud to Alibaba Cloud	Amazon RDS	Migrate data from an Amazon RDS for MySQL database to an ApsaraDB RDS for MySQL database	
		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 database	
	Amazon Aurora	Migrate data from an Amazon Aurora MySQL database to an ApsaraDB RDS for MySQL database	
		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 database	
	Amazon RDS for SQL Server	Migrate full data from an Amazon RDS for SQL Server database to an ApsaraDB RDS for SQL Server database	
	Migrate data between the same Alibaba Cloud account	RDS instance	Migrate data between RDS instances
		ApsaraDB RDS for MySQL instance	Migrate data from an ApsaraDB RDS for MySQL database to an ApsaraDB for PolarDB database
ApsaraDB RDS for MariaDB TX instance		Migrate data from an ApsaraDB RDS for MariaDB TX instance to an ApsaraDB RDS for MySQL instance	

Scenario	Source database	Topic
	Apsara PolarDB for MySQL cluster	Migrate data between ApsaraDB for PolarDB clusters
		Migrate data from an ApsaraDB for PolarDB database to an ApsaraDB RDS for MySQL database
	ApsaraDB for MongoDB instance	Migrate data from a standalone instance to a replica set instance or a sharded cluster instance
		Migrate data from a replica set instance to a sharded cluster instance
		Migrate the data of an ApsaraDB for MongoDB instance across regions
	Migrate data across Alibaba Cloud accounts	ApsaraDB for RDS instance
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
Migrate data from Alibaba Cloud to a user-created database	ApsaraDB RDS for MySQL instance	Migrate data from an ApsaraDB RDS for MySQL database to a user-created MySQL database
Migrate data between user-created databases	Oracle	Migrate data between user-created Oracle databases

2 Database and migration types

You can use DTS to migrate data between homogeneous and heterogeneous data sources. Typical scenarios include data migration to Alibaba Cloud, data migration between instances within Alibaba Cloud, and database splitting and scale-out. This topic describes the database types, database versions, and migration types that are supported by data migration.

For more information about data migration in various scenarios, see [Overview of data migration scenarios](#).

Migration types

Migration type	Description
Schema migration	<p>DTS migrates the schemas of the required objects from the source database to the destination database. Tables, views, triggers, and stored procedures can be migrated.</p> <p>For schema migration between heterogeneous databases, DTS converts the schema syntax based on the syntax of the source and destination databases. For example, it converts the NUMBER data type in Oracle databases to the DECIMAL data type in MySQL databases.</p>
Full data migration	<p>DTS migrates historical data of the required objects from the source database to the destination database. If you select only schema migration and full data migration, incremental data that is generated in the source database will not be migrated to the destination database.</p> <div data-bbox="427 1552 1434 1803"> Note: To ensure data consistency, do not write data into the source database during full data migration. To migrate data with minimal downtime, you must select schema migration, full data migration, and incremental data migration when configuring a data migration task.</div>

Migration type	Description
Incremental data migration	<p>DTS retrieves static snapshots that are used for full data migration from the source database and migrates the snapshot data to the destination database. Then, DTS synchronizes incremental data that is generated in the source database to the destination database in real time.</p> <div style="background-color: #f0f0f0; padding: 5px;">  Note: During incremental data migration, data between the source and destination databases is synchronized in real time. The migration task does not automatically stop. You must manually stop the migration task. </div>

Database and migration types

A user-created source or destination database, such as a MySQL, SQL Server, or Oracle database, can be one of the following types:

- **User-Created Database with Public IP Address**
- **Database without public IP:Port (Accessed through database gateway)**
- **Self built database accessed through Cloud Enterprise Network(CEN)**
- **User-Created Database in ECS Instance**
- **User-Created Database Connected Over Express Connect, VPN Gateway, or Smart Access Gateway**

Source database	Destination database	Migration type
<ul style="list-style-type: none"> • User-created MySQL database 5.1, 5.5, 5.6, 5.7, and 8.0 • ApsaraDB RDS for MySQL All versions 	User-created MySQL database 5.1, 5.5, 5.6, 5.7, and 8.0	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	ApsaraDB RDS for MySQL All versions	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration

Source database	Destination database	Migration type
	Apsara PolarDB for MySQL All versions	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	DRDS All versions <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: The database in DRDS instance are created based on ApsaraDB RDS for MySQL instances (MySQL 5.x) that you purchased. DTS does not support databases that are created based on ApsaraDB RDS for MySQL instances (MySQL 8.0) or ApsaraDB PolarDB for MySQL clusters. </div>	<ul style="list-style-type: none"> • Full data migration • Incremental data migration
	HybridDB for MySQL (Previous name : PetaData) All versions	<ul style="list-style-type: none"> • Full data migration • Incremental data migration
	AnalyticDB for MySQL 2.0 and 3.0	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	User-created PostgreSQL database 9.4, 9.5, 9.6, 10.x, 11.x, and 12	<ul style="list-style-type: none"> • Full data migration • Incremental data migration

Source database	Destination database	Migration type
	User-created Oracle database (RAC and non-RAC architecture) 9i, 10g, 11g, 12c, 18c, and 19c	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
ApsaraDB for MariaDB TX 10.3	ApsaraDB for MariaDB TX 10.3	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	ApsaraDB RDS for MySQL All versions	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	User-created MySQL database 5.1, 5.5, 5.6, 5.7, and 8.0	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
Apsara PolarDB for MySQL All versions	Apsara PolarDB for MySQL All versions	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration

Source database	Destination database	Migration type
	User-created MySQL database 5.1, 5.5, 5.6, 5.7, and 8.0	<ul style="list-style-type: none">• Schema migration• Full data migration• Incremental data migration
	ApsaraDB RDS for MySQL All versions	<ul style="list-style-type: none">• Schema migration• Full data migration• Incremental data migration
	AnalyticDB for MySQL 2.0 and 3.0	<ul style="list-style-type: none">• Schema migration• Full data migration• Incremental data migration

Source database	Destination database	Migration type
<ul style="list-style-type: none"> User-created SQL Server database 2005, 2008, 2008 R2, 2012, 2014, 2016, and 2017 <div data-bbox="268 501 699 927" style="background-color: #f2f2f2; padding: 5px;">  Note: <ul style="list-style-type: none"> SQL Server Cluster and SQL Server AlwaysOn High Availability Group are not supported. If the version of the source database is 2005, incremental data migration is not supported. </div> <ul style="list-style-type: none"> ApsaraDB RDS for SQL Server 2008, 2008 R2, 2012, 2014, 2016, and 2017 <div data-bbox="268 1120 699 1357" style="background-color: #f2f2f2; padding: 5px;">  Note: <p>If the version of the source database is 2008 or 2008 R2, incremental data migration is not supported.</p> </div>	<p>User-created SQL Server database 2005, 2008, 2008 R2, 2012, 2014, 2016, and 2017</p> <div data-bbox="724 470 1206 680" style="background-color: #f2f2f2; padding: 5px;">  Note: <p>SQL Server Cluster and SQL Server AlwaysOn High Availability Group are not supported.</p> </div> <p>ApsaraDB RDS for SQL Server 2008, 2008 R2, 2012, 2014, 2016, and 2017</p>	<ul style="list-style-type: none"> Schema migration Full data migration Incremental data migration <ul style="list-style-type: none"> Schema migration Full data migration Incremental data migration
<p>User-created Oracle database (RAC and non-RAC architecture) 9i, 10g, 11g, 12c, 18c, and 19c</p>	<p>User-created Oracle database (RAC and non-RAC architecture) 9i, 10g, 11g, 12c, 18c, and 19c</p> <p>PolarDB database compatible with Oracle All versions</p>	<ul style="list-style-type: none"> Schema migration Full data migration Incremental data migration <ul style="list-style-type: none"> Schema migration Full data migration Incremental data migration

Source database	Destination database	Migration type
	ApsaraDB RDS for PPAS 9.3 and 10	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	User-created MySQL database 5.1, 5.5, 5.6, 5.7, and 8.0	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	ApsaraDB RDS for MySQL All versions	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
	DRDS All versions <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: The database in DRDS instance are created based on ApsaraDB RDS for MySQL instances (MySQL 5.x) that you purchased. DTS does not support databases that are created based on ApsaraDB RDS for MySQL instances (MySQL 8.0) or ApsaraDB PolarDB for MySQL clusters. </div>	<ul style="list-style-type: none"> • Full data migration • Incremental data migration

Source database	Destination database	Migration type
	AnalyticDB for MySQL 2.0 and 3.0	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
User-created PostgreSQL database 9.4, 9.5, 9.6, 10.x, 11.x, and 12	<ul style="list-style-type: none"> • User-created PostgreSQL database 9.4, 9.5, 9.6, 10.x, 11.x, and 12 • ApsaraDB RDS for PostgreSQL 9.4, 10, 11, and 12 	<ul style="list-style-type: none"> • Schema migration • Full data migration • Incremental data migration
<ul style="list-style-type: none"> • User-created MongoDB database (single-node, replica set, and sharded cluster architecture) 3.0, 3.2, 3.4, 3.6, and 4.0 • ApsaraDB for MongoDB instance (single-node and replica set architecture) All versions 	<ul style="list-style-type: none"> • User-created MongoDB database (single-node, replica set, and sharded cluster architecture) 3.0, 3.2, 3.4, 3.6, and 4.0 • ApsaraDB for MongoDB instance (single-node, replica set, and sharded cluster architecture) All versions 	<ul style="list-style-type: none"> • Full data migration • Incremental data migration <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 5px; margin-top: 10px;">  Note: MongoDB is a NoSQL database that does not require schema migration. </div>

Source database	Destination database	Migration type
User-created Redis database (single-host architecture) 2.8, 3.0, 3.2, 4.0, and 5.0	User-created Redis database (single-host and cluster architecture) 2.8, 3.0, 3.2, 4.0, and 5.0	<ul style="list-style-type: none"> Full data migration Incremental data migration <div style="border: 1px solid gray; background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: Redis is a NoSQL database that does not require schema migration. </div>
	ApsaraDB for Redis instance (single-host and cluster architecture) Community versions 4.0 and 5.0	<ul style="list-style-type: none"> Full data migration Incremental data migration
User-created TiDB database	User-created MySQL database 5.1, 5.5, 5.6, 5.7, and 8.0	<ul style="list-style-type: none"> Schema migration Full data migration Incremental data migration
	ApsaraDB RDS for MySQL All versions	<ul style="list-style-type: none"> Schema migration Full data migration Incremental data migration

Source database	Destination database	Migration type
	Apsara PolarDB for MySQL All versions	<ul style="list-style-type: none">• Schema migration• Full data migration• Incremental data migration
User-created Db2 database 9.7 and 10.5	User-created MySQL database 5.1, 5.5, 5.6, 5.7, and 8.0	<ul style="list-style-type: none">• Schema migration• Full data migration• Incremental data migration

3 Data type mappings between heterogeneous databases

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. This topic lists the data type mappings for you to evaluate the impact of data migration on your business.

For more information about data migration scenarios, see [Overview of data migration scenarios](#).

Migrate data from a user-created TiDB database to a user-created MySQL database, an ApsaraDB RDS for MySQL database, or an ApsaraDB for POLARDB database

Data type in the TiDB database	Data type in the MySQL database
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	INTEGER
JSON	JSON
MEDIUMBLOB/LONGBLOB	MEDIUMBLOB/LONGBLOB
TINYBLOB/BLOB/	TINYBLOB/BLOB/
MEDIUMINT	MEDIUMINT
SET	SET

Data type in the TiDB database	Data type in the MySQL database
SMALLINT	SMALLINT
TEXT/LONGTEXT	TEXT/LONGTEXT
TIME	TIME
TIMESTAMP	TIMESTAMP
TINYINT	TINYINT
VARBINARY	VARBINARY
VARCHAR	VARCHAR
YEAR	YEAR

Migrate data from a user-created Oracle database to a user-created MySQL database, an ApsaraDB RDS for MySQL database, or an ApsaraDB for POLARDB database

Data type in the Oracle database	Data type in the MySQL database	Supported by DTS
varchar2(n [char/byte])	varchar(n)	Yes
nvarchar2[(n)]	national varchar[(n)]	Yes
char[(n [byte/char])]	char[(n)]	Yes
nchar[(n)]	national char[(n)]	Yes
number[(p[,s])]	decimal[(p[,s])]	Yes
float(p)	double	Yes
long	longtext	Yes
date	datetime	Yes
binary_float	decimal(65,8)	Yes
binary_double	double	Yes
timestamp[(fractional_seconds_precision)]	datetime[(fractional_seconds_precision)]	Yes
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]	Yes
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]	Yes
clob	longtext	Yes
nclob	longtext	Yes

Data type in the Oracle database	Data type in the MySQL database	Supported by DTS
blob	longblob	Yes
raw	varbinary(2000)	Yes
long raw	longblob	Yes
bfile	N/A	No
interval year(year_precision) to month	N/A	No
interval day(day_precision)to second[(fractional_seconds_precision)]	N/A	No

**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 these three data types . You must make sure that columns with these three data types are excluded from the objects to be migrated.

- The timestamp data type of MySQL databases does not contain the time zone information. However, the timestamp with time zone and timestamp with local time zone data types in Oracle databases provide time zone information. Therefore, DTS converts the values of these data types based on the time zone to UTC time for storage in the destination instance.

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

Data type in the Oracle database	Data type in the DRDS database	Supported by DTS
varchar2(n [char/byte])	varchar(n)	Yes
nvarchar2[(n)]	national varchar[(n)]	Yes
char[(n [byte/char])]	char[(n)]	Yes
nchar[(n)]	national char[(n)]	Yes
number[(p[,s])]	decimal[(p[,s])]	Yes
float(p)	double	Yes

Data type in the Oracle database	Data type in the DRDS database	Supported by DTS
long	longtext	Yes
date	datetime	Yes
binary_float	decimal(65,8)	Yes
binary_double	double	Yes
timestamp[(fractional_seconds_precision)]	datetime[(fractional_seconds_precision)]	Yes
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]	Yes
timestamp[(fractional_seconds_precision)]with localtimezone	datetime[(fractional_seconds_precision)]	Yes
clob	longtext	Yes
nclob	longtext	Yes
blob	longblob	Yes
raw	varbinary(2000)	Yes
long raw	longblob	Yes
bfile	N/A	No
interval year(year_precision) to month	N/A	No
interval day(day_precision)to second[(fractional_seconds_precision)]	N/A	No

**Note:**

- A char column with a length greater than 255 bytes is converted to the varchar(n) type in the DRDS database.
- The timestamp data type of DRDS databases does not contain the time zone information. However, the timestamp with time zone and timestamp with local time zone data types in Oracle databases provide time zone information. Therefore, DTS converts the values of these data types based on the time zone to UTC time for storage in the destination instance.

Migrate data from a user-created Oracle database to an ApsaraDB RDS for PPAS database

Data type in the Oracle database	Data type in the ApsaraDB RDS for PPAS database	Supported by DTS
varchar2(n [char/byte])	varchar2[(n)]	Yes
nvarchar2[(n)]	nvarchar2[(n)]	Yes
char[(n [byte/char])]	char[(n)]	Yes
nchar[(n)]	nchar[(n)]	Yes
number[(p[,s])]	number[(p[,s])]	Yes
float(p)	double precision	Yes
long	long	Yes
date	date	Yes
binary_float	real	Yes
binary_double	double precision	Yes
timestamp[(fractional_seconds_precision)]	timestamp[(fractional_seconds_precision)]	Yes
timestamp[(fractional_seconds_precision)]with time zone	timestamp[(fractional_seconds_precision)]with time zone	Yes
timestamp[(fractional_seconds_precision)]with local time zone	timestamp[(fractional_seconds_precision)]with time zone	Yes
clob	clob	Yes
nclob	nclob	Yes
blob	blob	Yes
raw	raw(size)	Yes
long raw	long raw	Yes
bfile	N/A	No
interval year(year_precision) to month	interval year to month	No
interval day(day_precision) to second[(fractional_seconds_precision)]	interval day to second[(fractional_seconds_precision)]	No

**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 converts the values of these data types based on the time zone to UTC time before storing them in the timestamp[(fractional_seconds_precision)]with time zone data type in the destination instance.

4 Migration task management

4.1 Object name mapping

DTS provides the object name mapping feature. You can use this feature to change the name of an object in the destination instance. This topic describes how to use the object name mapping feature when you configure a data migration task.

Limits

- While configuring a data migration task, you can use the object name mapping feature when you go to the step of **Configure Migration Types and Objects**.

**Note:**

Don't use the object name mapping feature once a data migration task started. Otherwise, the migration might fail.

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

Non-batch mappings of databases, tables, and columns

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

**Note:**

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

1.Configure Source and Destination 2.Configure Migration Types and Objects 3.Map name modification 4.Precheck

* Migration Types: Schema Migration Full Data Migration Incremental Data Migration

Data migration applies to short-term migration scenarios. Typical scenarios include migrating data to the cloud, scaling and sharding databases, and migrating data between Apsara Stack databases.
For long-term data synchronization in real time, use the data synchronization feature.

Available

If you search globally, please expand the:

- dtstestdata
 - Tables
 - Views
 - Functions
 - Procedures
 - sys

Select All

Selected (To edit an object name or its filter, hover over the object and click Edit.) [Learn more.](#)

- dtstestdata (2Objects)
 - customer
 - order Edit

Remove All

*Name batch change: No Yes

Information:
1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.
2. DDL operations are not supported during data migration because this can cause migration failures.

Cancel Previous Save Precheck

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



Note:

- DDL and DML statement filtering functions are also supported when performing database or table name mapping, when the source and destination database are user-create MySQL, ApsaraDB RDS for MySQL, or ApsaraDB PolarDB for MySQL. You can select DDL or DML statements that need to be synchronized according to business requirements.

- If the selected statements of the database name mapping and table name mapping are different, the statement selected during table name mapping shall prevail.

- Database name mapping

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

Information: After you edit the source database name, the name of the destination database is also updated.

* DatabaseName: dtstestdatanew

Source Database Name:dtstestdata

OK

- Table name mapping

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

Edit Table

Information: After you edit the table or column name in the source database, the corresponding table or column name in the destination database is also updated.

Source Table Name: customer

* Table Name:

Filter:

<input checked="" type="checkbox"/> Select All	Column Name	Type
<input checked="" type="checkbox"/>	<input type="text" value="address"/>	varchar(32)
<input checked="" type="checkbox"/>	<input type="text" value="id"/>	int(11)
<input checked="" type="checkbox"/>	<input type="text" value="name"/>	varchar(32)

- Column name mapping

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

Edit Table
✕

Information: After you edit the table or column name in the source database, the corresponding table or column name in the destination database is also updated.

* Table Name:

Filter:

DTS supports the WHERE clause in SQL statements. Only data that meets the WHERE clause can be migrated to the destination

Verify

Select

All

Column Name

Source Column Name:address

<input checked="" type="checkbox"/>	<input style="border: 1px solid green;" type="text" value="addressnew"/>		varchar(32)
<input checked="" type="checkbox"/>	<input type="text" value="id"/>		int(11)
<input checked="" type="checkbox"/>	<input type="text" value="name"/>		varchar(32)

OK



Note:

In this step, you can clear the options of columns that do not need to be synchronized.

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

Batch mappings of tables, and columns

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



Notice:

If the granularity of the last selected object to be migrated is database, the **Name batch change** Option is not selectable.

2. At the bottom of the page, click **Yes**, and then click the **Advanced Settings**.

The screenshot displays the '2.Configure Migration Types and Objects' step of a migration wizard. At the top, there are four tabs: '1.Configure Source and Destination Databases', '2.Configure Migration Types and Objects' (active), '3.Advanced Settings', and '4.Precheck'. Below the tabs, there are migration type options: 'Schema Migration' (checked), 'Full Data Migration' (checked), and 'Incremental Data Migration' (unchecked). A note states: 'During full data migration, data updates in the source database are not migrated to the destination instance. For data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration.' The main area is split into 'Available' and 'Selected' sections. The 'Available' section shows a tree view of databases and their objects (Tables, Views, Functions, Procedures). A red circle '1' highlights the 'Yes' radio button for 'Name batch change'. The 'Selected' section shows a list of selected objects: 'dtstestdata (2Objects)', 'customer', and 'order'. A red circle '2' highlights the 'Advanced Settings' button at the bottom right of the interface.

3. In the **Advanced Settings** page, batch map table and column names as needed.

a. Select the choice range on your requirements. In this case, Select **Section**.



Notice:

If you select **Section**, you can enter the database name or table name keyword to filter the databases or tables to be mapped.

b. Select a rule and configure the rule details as required.

c. Click **Add Modification Rules**.



Notice:

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

- d. Click **Preview** to check whether the rules you set meet expectations.

- e. In the **Partial Preview** tab, click the  icon at the end of the database name listed in **Destination Database** box ,and then click **Column Name Modification**.
- f. Configure mapping rules for column names as required. After the configuration, click **OK**.

Column selection and column name changes
✕

customer table 3 columns

Filter: Verify

1

Search

2

Name Added Suffix ▾

Prefix:

Suffix:

3

OK

Modify Rules 1: Table Namecustomer_Column NameAdd Prefix: test_,Suffix:
✕

Original list name	Type	Target column name
<input type="text" value="address"/>	varchar(32)	<input type="text" value="test_address"/>
<input type="text" value="id"/>	int(11)	<input type="text" value="test_id"/>
<input type="text" value="name"/>	varchar(32)	<input type="text" value="test_name"/>

Cancel
OK



Notice:

- The default rule applies to all columns, you can enter a keyword of the column name, and then click **Search** to filter the columns whose names are mapped.
- In the dialog box, you can also configure filters to filter the data to be migrated. For more information, see [filter configuration](#).

g. Click **OK**.

4. Click **Precheck**.

5. Configure other parameters that are required for the data migration task.

4.2 Scheduling policies of regular data migration

Scheduling is an advanced feature of data migration. You can configure a scheduling policy and scheduling cycle when you migrate schemas and historical data from the source database to the destination database on a regular basis. The scheduling feature allows you to flexibly build data warehouses. For example, you can build a T+1 data warehouse. This topic describes the scheduling policies of regular data migration, the limits on scheduling policies, and the scenarios of using these policies.

Supported databases

All databases that support full data migration support the scheduling feature. For more information, see [Database and migration types](#).

Supported regions

China (Qingdao), China (Beijing), China (Zhangjiakou-Beijing Winter Olympics), China (Shenzhen), and China (Shanghai) are supported.

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 #unique_51 .

Configuration method

You can configure a scheduling policy for a regular data migration task in the **Configure Migration Types and Objects** step.

**Note:**

For more information about how to configure a scheduling policy, see [Migrate full data between ApsaraDB RDS for MySQL instances on a regular basis](#).

Scheduling policies

Scheduling policy	Description	Scenario
Duplicate mode	<p>DTS migrates schemas and historical data of the required objects from the source database to the destination database based on the scheduling cycle.</p> <div style="background-color: #f0f0f0; padding: 10px;">  <p>Note:</p> <ul style="list-style-type: none"> • During task scheduling, DTS automatically creates a database in the destination instance to receive migrated data. The database name is suffixed with <code>_yyyymmdd_HH_mm_ss</code>, which indicates the scheduled time. For example, a data migration task migrates data from a database named <code>dtstestdata</code>. The task is scheduled at 11:00:00 on February 10, 2020. When DTS migrates data at this time, it creates a database named <code>dtstestdata_20200207_16_00_00</code> in the destination instance. • The destination database must have sufficient storage space to store the data that is migrated during each scheduling cycle. </div>	<p>This policy is applicable to test scenarios where test operations are performed on multiple replicas of historical data.</p>
Resynchronize mode	<p>At the first scheduled time, DTS migrates schemas and historical data of the required objects to the destination database. At each subsequent scheduled time, DTS migrates schemas and historical data of the required objects to the destination database again.</p> <div style="background-color: #f0f0f0; padding: 10px;">  <p>Warning:</p> <ul style="list-style-type: none"> • Before each scheduled time, you must manually delete the schemas and data that have been migrated to the destination database at the last scheduled time. Otherwise, the data migration fails. • If DTS does not support schema migration from a specific database, you must create a schema in the destination database based on the schema of the required objects. </div>	<p>This policy is applicable to scenarios where only the latest historical data is maintained in the destination database for testing or data analysis.</p>

Scheduling policy	Description	Scenario
Based on the incremental mode of time stamp	<p>DTS migrates incremental data that is generated after the previous scheduling to the destination database based on the specified timestamp. This policy has the following limits:</p> <ul style="list-style-type: none"> • This policy is supported only when the source and destination databases are MySQL databases. • Each source table must contain a timestamp field. You must specify the timestamp field of each table when you configure a data migration task. • Source tables cannot have FOREIGN KEY constraints. Otherwise, data migration may fail. If you update a child table that has FOREIGN KEY constraints in the destination database and then update the attached table, an error occurs when data is inserted into the child table. • Data that is deleted by using the DELETE or TRUNCATE command cannot be migrated. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> • If a table does not have a primary key, DTS uses an entire data record as the primary key. This means that a data record is matched only if all fields in the data record are matched. DTS deletes matched data records and inserts new data records. • If a primary key conflict occurs when a data record is written to the destination database, DTS runs the REPLACE INTO command to insert the data record. </div>	<p>This policy is applicable to lightweight data migration scenarios where you do not need to migrate all historical data each time, such as building a T+1 data warehouse.</p>

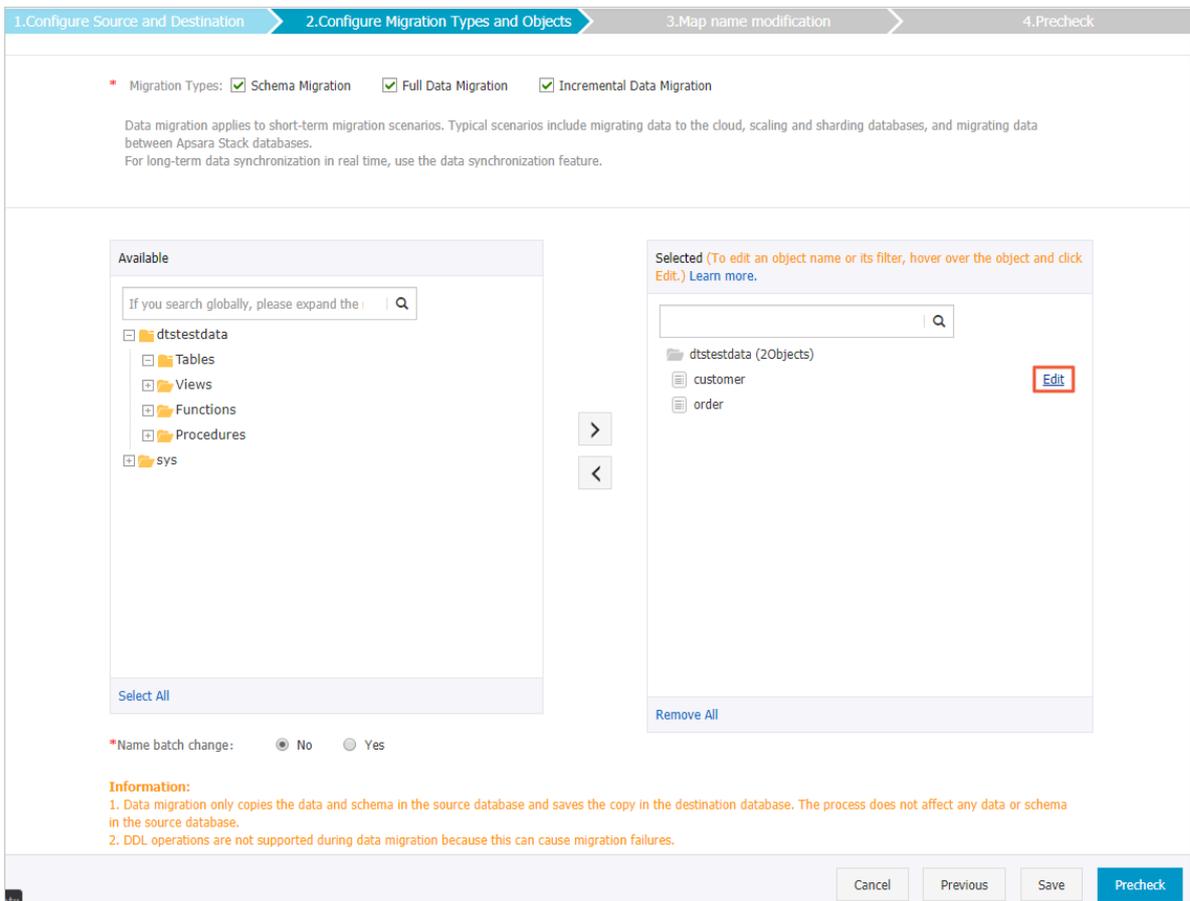
4.3 Use conditions to filter data

When configuring the objects to be migrated in a data migration task, you can specify conditions to filter data. Only data that meets the specified conditions is migrated to the

destination database. This feature is applicable to scenarios such as regular data migration and table partitioning.

Procedure

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



2. In the Edit Table dialog box that appears, enter an SQL condition in the **Filter** field.

- Relational databases, such as MySQL and SQL Server.

✕

Information: After you edit the table or column name in the source database, the corresponding table or column name in the destination database is also updated.

* Table Name:

Filter:

<input checked="" type="checkbox"/> Select All	Column Name	Type
<input checked="" type="checkbox"/>	<input type="text" value="address"/>	varchar(32)
<input checked="" type="checkbox"/>	<input type="text" value="orderid"/>	int(11)
<input checked="" type="checkbox"/>	<input type="text" value="name"/>	varchar(32)



Note:

- An SQL condition is a standard SQL WHERE statement. The following operators are supported: =, !=, <, and >. Only data that meets the WHERE condition is migrated to the destination database. In this example, enter `orderid>100`.

- You can use apostrophes (') in an SQL condition if necessary. For example, you can enter `address in('hangzhou','shanghai')`.

- Non-relational databases, such as MongoDB.

Edit Collection Name
✕

Information: If a collection name in the source database is edited, the corresponding collection name in the destination database is also updated.

* Collection Name:

Filter:

```
{ "id": { "$gt": "52b64e55dc4449b7ba3d1183a0ea171b" } }
```

Verify

OK



Note:

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

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



Note:

- If the syntax is valid, the **Information** message prompts that **the validation is passed**.
- If the syntax is invalid, the **Error** message appears, you must modify the condition based on the instructions.

4. Click **OK**.

5. Configure other parameters that are required for the data migration task.

4.4 View the progress of a data migration task

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

Prerequisites

A data migration task is started.

Procedure

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



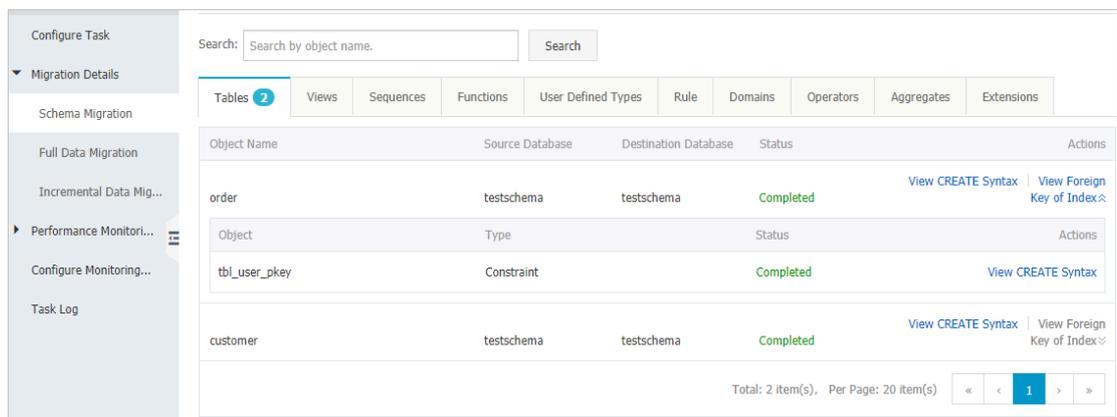
4. On the **Migration Tasks** page, click the ID of the destination instance.
5. Perform operations based on the migration type.



Note:

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

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



The screenshot shows the 'Migration Details' section of the interface. On the left, a navigation pane lists options: 'Configure Task', 'Migration Details' (expanded), 'Schema Migration', 'Full Data Migration', 'Incremental Data Mig...', 'Performance Monitori...', 'Configure Monitoring...', and 'Task Log'. The main area has a search bar and tabs for 'Tables' (selected), 'Views', 'Sequences', 'Functions', 'User Defined Types', 'Rule', 'Domains', 'Operators', 'Aggregates', and 'Extensions'. Below the tabs is a table with columns: Object Name, Source Database, Destination Database, Status, and Actions. The table lists three objects: 'order', 'tbl_user_pkey', and 'customer'. The 'order' object is expanded to show a sub-table with columns: Object, Type, Status, and Actions. The sub-table lists 'tbl_user_pkey' as a 'Constraint' with a 'Completed' status. The main table also shows 'customer' with a 'Completed' status. At the bottom, there is a pagination bar indicating 'Total: 2 item(s), Per Page: 20 item(s)' and navigation arrows.

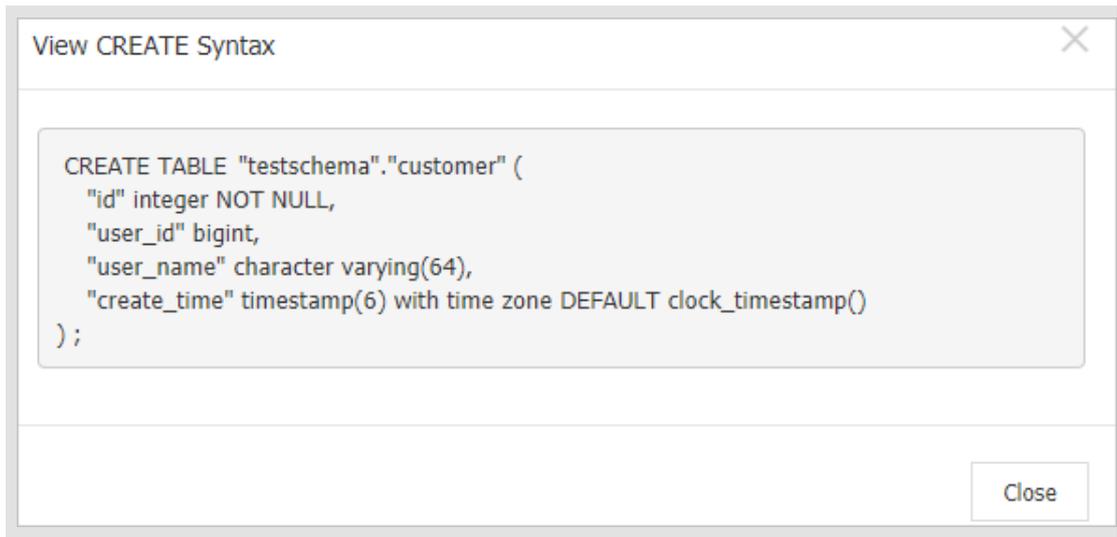
Object Name	Source Database	Destination Database	Status	Actions								
order	testschema	testschema	Completed	View CREATE Syntax View Foreign Key of Index								
<table border="1"> <thead> <tr> <th>Object</th> <th>Type</th> <th>Status</th> <th>Actions</th> </tr> </thead> <tbody> <tr> <td>tbl_user_pkey</td> <td>Constraint</td> <td>Completed</td> <td>View CREATE Syntax</td> </tr> </tbody> </table>					Object	Type	Status	Actions	tbl_user_pkey	Constraint	Completed	View CREATE Syntax
Object	Type	Status	Actions									
tbl_user_pkey	Constraint	Completed	View CREATE Syntax									
customer	testschema	testschema	Completed	View CREATE Syntax View Foreign Key of Index								

Total: 2 item(s), Per Page: 20 item(s)

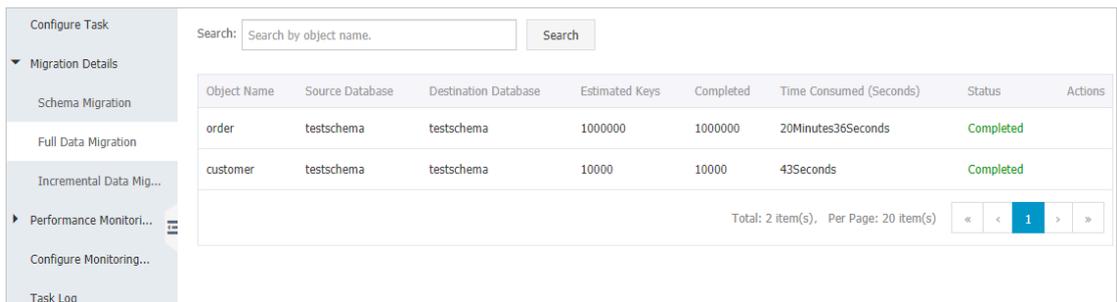


Note:

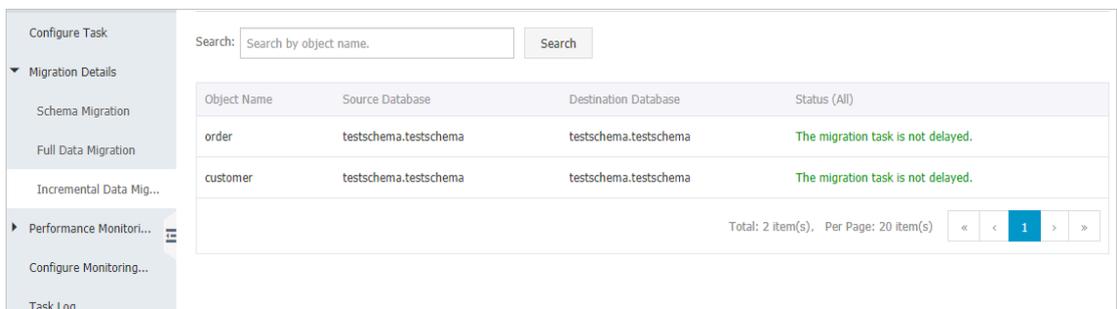
In the Actions column of an object, click **View Creation Syntax** or choose **View Index Foreign Keys > View Creation Syntax**. In the View Creation Syntax message that appears, you can view the creation syntax of the object.



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



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



Reference

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

4.5 View the performance of full data migration

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

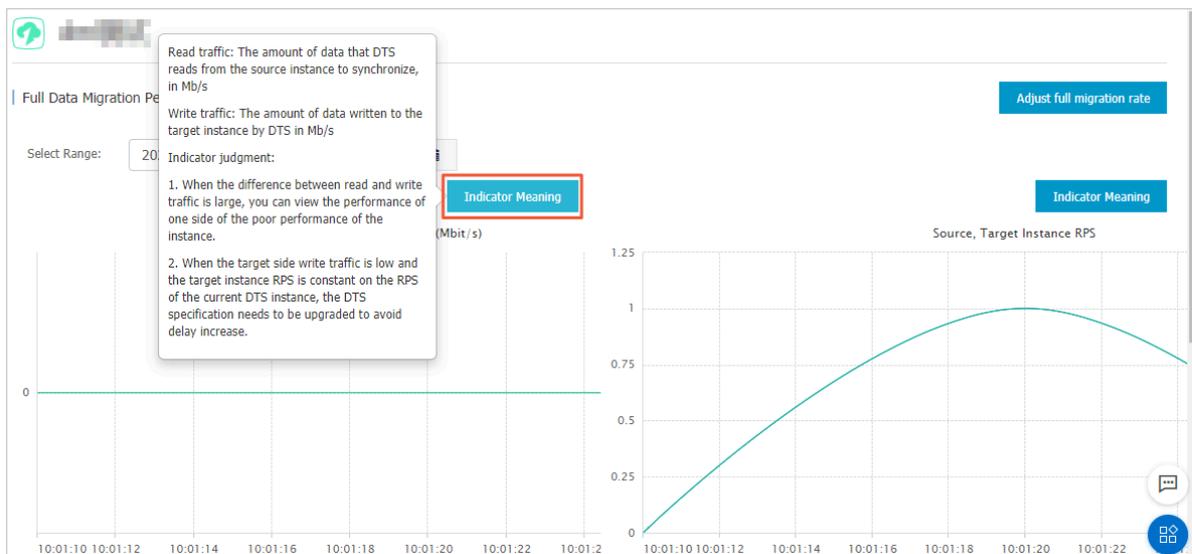
Procedure

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



4. On the **Migration Tasks** page, click the ID of the destination instance.
5. In the left-side navigation pane, click **Performance Monitoring > Full Data Migration Performance**.

6. Select a time range to view the trend charts for the performance of full data migration.



Note:

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

4.6 Check the performance of incremental data migration

DTS provides three performance metrics: migrated rows, bandwidth, and migration performance (QPS). You can monitor the progress of incremental data migration tasks in the console by using the performance metrics.

Procedure

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



4. On the **Migration Tasks** page, click the ID of the destination 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.

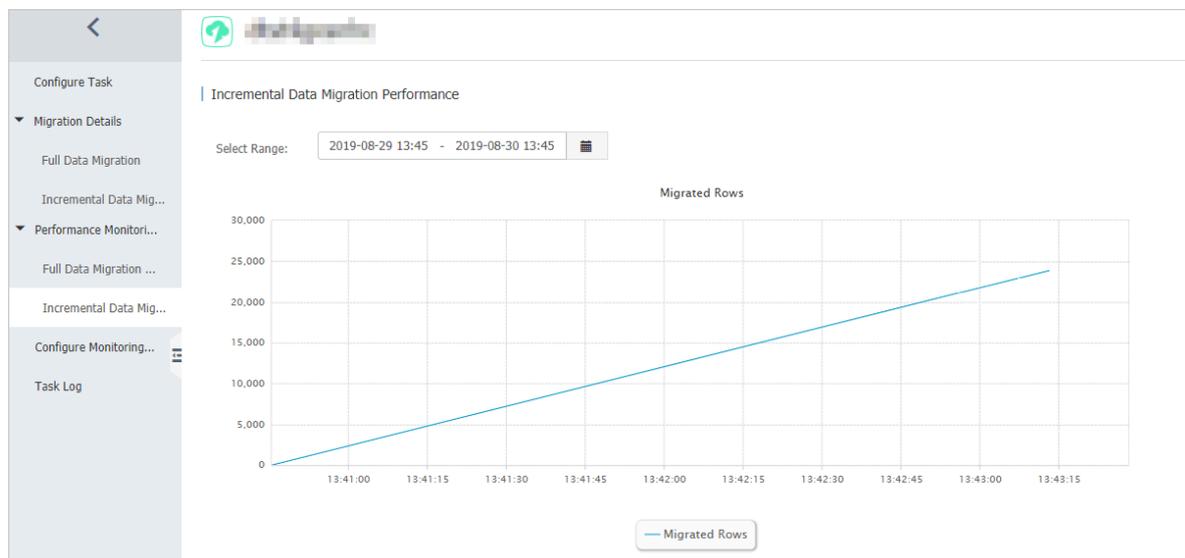


Table 4-1: Performance metrics

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.

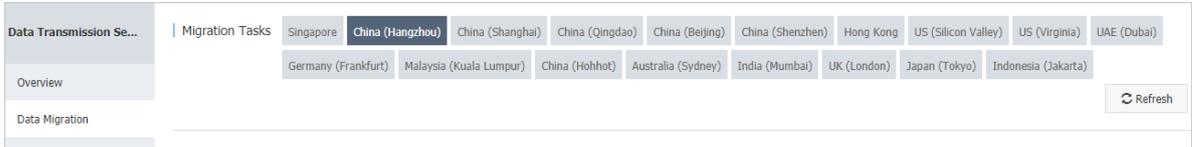
4.7 Adjust the full migration rate

During data migration, you can adjust the full migration rate according to the read and write pressure of the source database.

Procedure

1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of the **Migration Tasks** page, select the region where the destination instance resides.

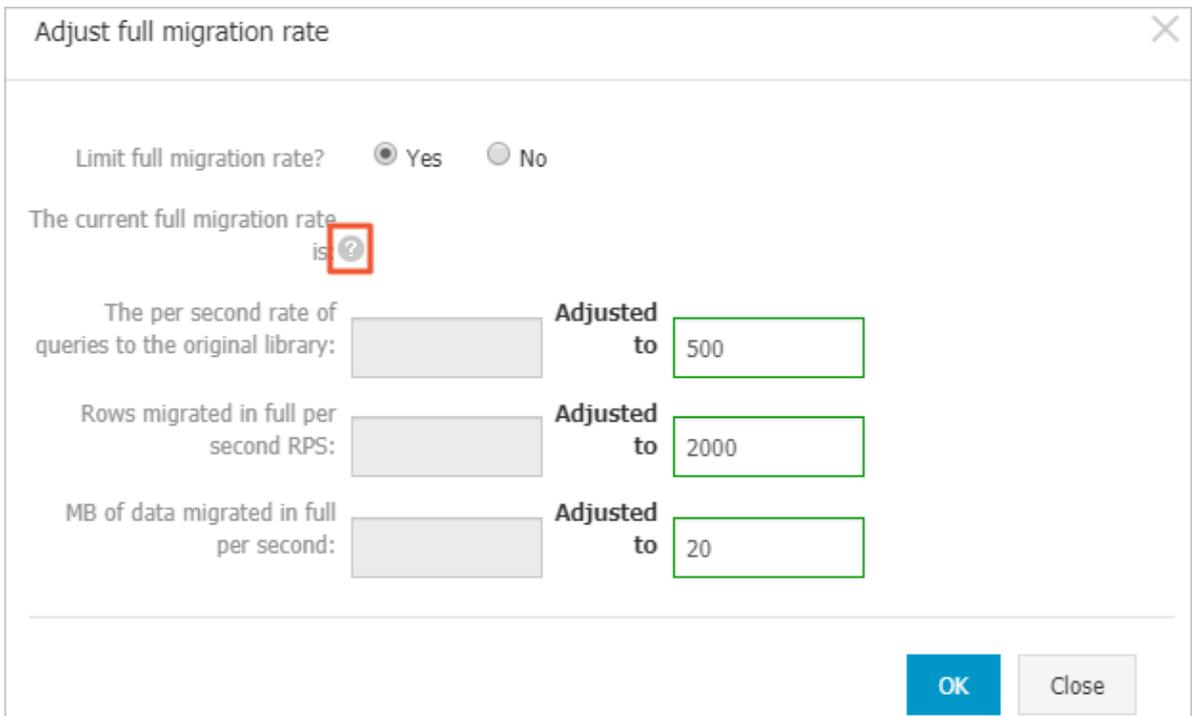


4. On the **Migration Tasks** page, click the ID of the destination instance.

5. In the left-side navigation pane, click **Performance Monitoring > Full Data Migration Performance**.

6. Click **adjust the full migration rate**.

7. Specify whether to limit the full migration rate and adjust the related parameters according to business requirements.



Note:

You can click  icon to see a detailed explanation of each parameter.

8. Click **OK**.

4.8 Fix a failed data migration task

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

Fix a failed task during schema migration

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

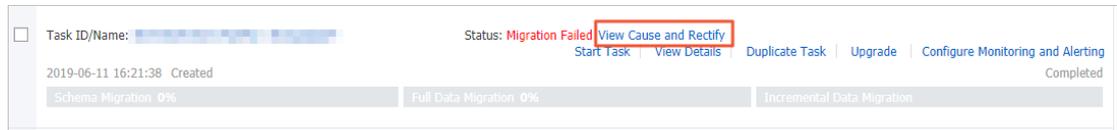
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of the **Migration Tasks** page, select the region where the destination instance resides.



4. Use one of the following methods to fix the failed task:

- Method 1

a. Find the target task and click **View Cause and Rectify**.



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

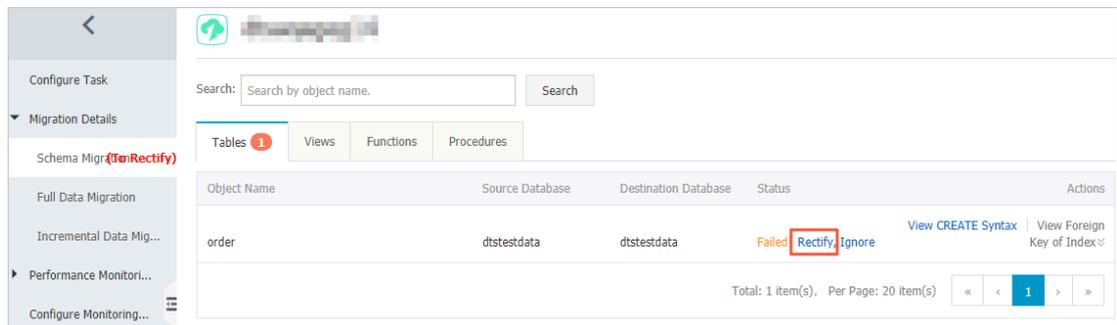
c. Click **Restart Task**.

- Method 2

a. Click the ID of the target task.

b. In the left-side navigation pane, choose **Migration Details > Schema Migration**.

c. On the **Schema Migration** page, find the object that causes the migration failure and click **Rectify** in the Status column.



d. Fix the issue based on the cause of failure that is displayed in the **Rectify** dialog box. For example, you can fix an issue by modifying the schema syntax.

e. Click **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 the troubleshooting is successful.
- If the troubleshooting is successful, the **Schema Migration** page appears and the status of the object changes to **Finished**.

5. If no objects are in the Failed state, DTS proceeds with the data migration task, for example, entering the full data migration process.

Fix a failed task during full data migration

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



4. Find the target task and click **View Cause and Rectify**.

DTS allows you to fix a task that fails during full data migration due to the following reasons.



Note:

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

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

4.10 Stop a data migration task

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

Prerequisites

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

Impact

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

Procedure

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



4. Perform operations based on the selected migration type.



Warning:

To minimize the negative impact of data migration on your business, we recommend that you switch your workloads to the destination instance and prepare a rollback solution. For more information, see [#unique_59](#). If you do not need to switch your workloads, you can stop the migration task by using the following procedure.

- Full data migration

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

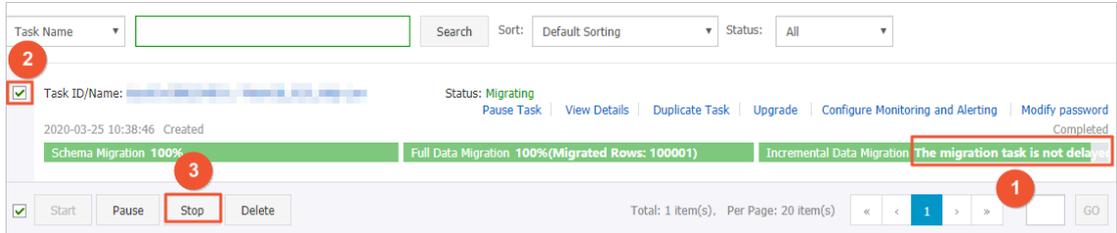
- Incremental data migration

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

- a. Wait until the task progress bar shows **Incremental Data Migration** and **The migration task is not delayed**. Then, stop writing data to the source database for a

few minutes. In some cases, the progress bar shows the delay time of **incremental data migration**.

- b. After the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



5 Precheck and troubleshooting

5.1 Source database connectivity

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

The following shows the potential causes of the failed source database connectivity check.

The database account or password is incorrect

Troubleshooting:

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

**Note:**

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

Solution:

Log on to the [DTS console](#), enter the correct account and password, and perform another precheck.

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

Troubleshooting:

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

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

**Note:**

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

Check whether the authorized IP address list includes the IP addresses of the DTS servers. For more information, see [#unique_62](#).

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

Solution:

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

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



Note:

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

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

Log on to the [DTS console](#) to perform another precheck.

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

Troubleshooting:

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

Solution:

After you disable the firewall, log on to the [DTS console](#) to perform another precheck.

Network connections fail

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

5.2 Source database version

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

DTS checks whether the source database version meets the requirements. For more information, see [Database and migration types](#).

6 Migration from a user-created database to Alibaba Cloud

6.1 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 [#unique_65](#).
- 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

- Schema migration

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



Note:

- During schema migration, DTS changes the value of the SECURITY attribute 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.

- Incremental data 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
DDL	<ul style="list-style-type: none"> • 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

Database	Schema migration	Full data migration	Incremental data migration
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:

- [#unique_66](#) 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

[#unique_66](#)

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.

The screenshot shows the '1. Configure Source and Destination' step of the DTS migration task. The 'Task Name' is 'MySQL_TO_MySQL'. The 'Source Database' section is configured with: Instance Type: 'User-Created Database with Public IP Address', Instance Region: 'Singapore', Database Type: 'MySQL', Hostname or IP Address: [redacted], Port Number: '3306', Database Account: 'dtstest', and Database Password: [redacted]. A 'Test Connectivity' button shows 'Passed'. The 'Destination Database' section is configured with: Instance Type: 'RDS Instance', Instance Region: 'Singapore', RDS Instance ID: [redacted], Database Account: 'dtstest', and Database Password: [redacted]. A 'Test Connectivity' button also shows 'Passed'. At the bottom right, there are 'Cancel' and 'Set Whitelist and Next' buttons.

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.
Source Database	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. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.</p> <p> 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.</p>
	Database Type	Select MySQL .
	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 .
	Database 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 .
	Database Password	<p>Enter the password for the source database account.</p> <p> 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.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination RDS instance resides.
	RDS Instance ID	Select the ID of the destination RDS instance.

Section	Parameter	Description
	Database 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 .
	Database Password	Enter the password for the destination database account.  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.
	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.

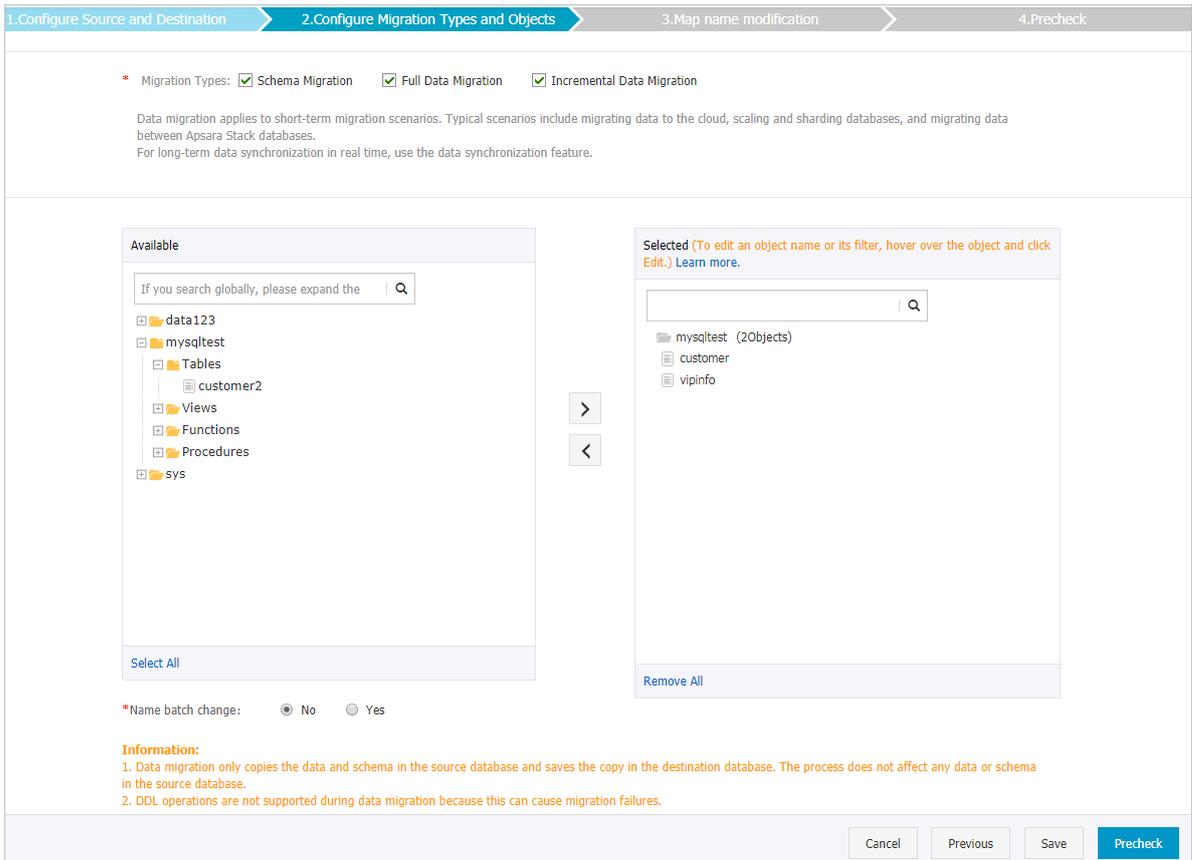
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <p>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.</p>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> 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. </div>

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



Note:

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

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 [#unique_59](#). If you do not need to switch your workloads, you can stop the migration task by using the following procedure.

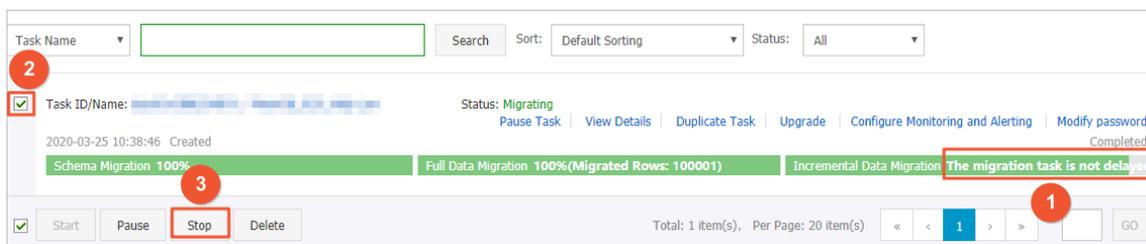
- Full data migration

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

- Incremental data migration

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

1. 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**.
2. 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.

6.2 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 [#unique_68](#).

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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

- Schema migration

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



Note:

- During schema migration, DTS changes the value of the SECURITY attribute 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.

- Incremental data 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
DDL	<ul style="list-style-type: none"> • 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 more information about how to create and authorize a database account, see the following topics:

- [#unique_66](#) 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

1. [#unique_66](#).
2. [#unique_69](#).

Procedure

1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the 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.

The screenshot shows the '1. Configure Source and Destination Databases' step in the DTS console. It includes fields for Task Name, Source Database (Instance Type, Region, Peer VPC, Database Type, IP Address, Port Number, Database Account, Password), and Destination Database (Instance Type, Region, RDS Instance ID, Database Account, Password, Encryption). Buttons for 'Test Connectivity' are present for both source and destination.

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 Database	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 .

Section	Parameter	Description
	Database 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 .
	Database Password	Enter the password for the database account.  Note: After the source database parameters are specified, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are correct. If the source database parameters are correct, the Test Passed message is displayed, If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database parameters as prompted.
Destination Database	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 .
	Database Password	Enter the password for the database account.  Note: After the destination database parameters are specified, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are correct. If the destination database parameters are correct, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database parameters as prompted.

Section	Parameter	Description
	Encryption	<p>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.</p> <p> Note: The Encryption parameter is available only in mainland China and Hong Kong(China).</p>

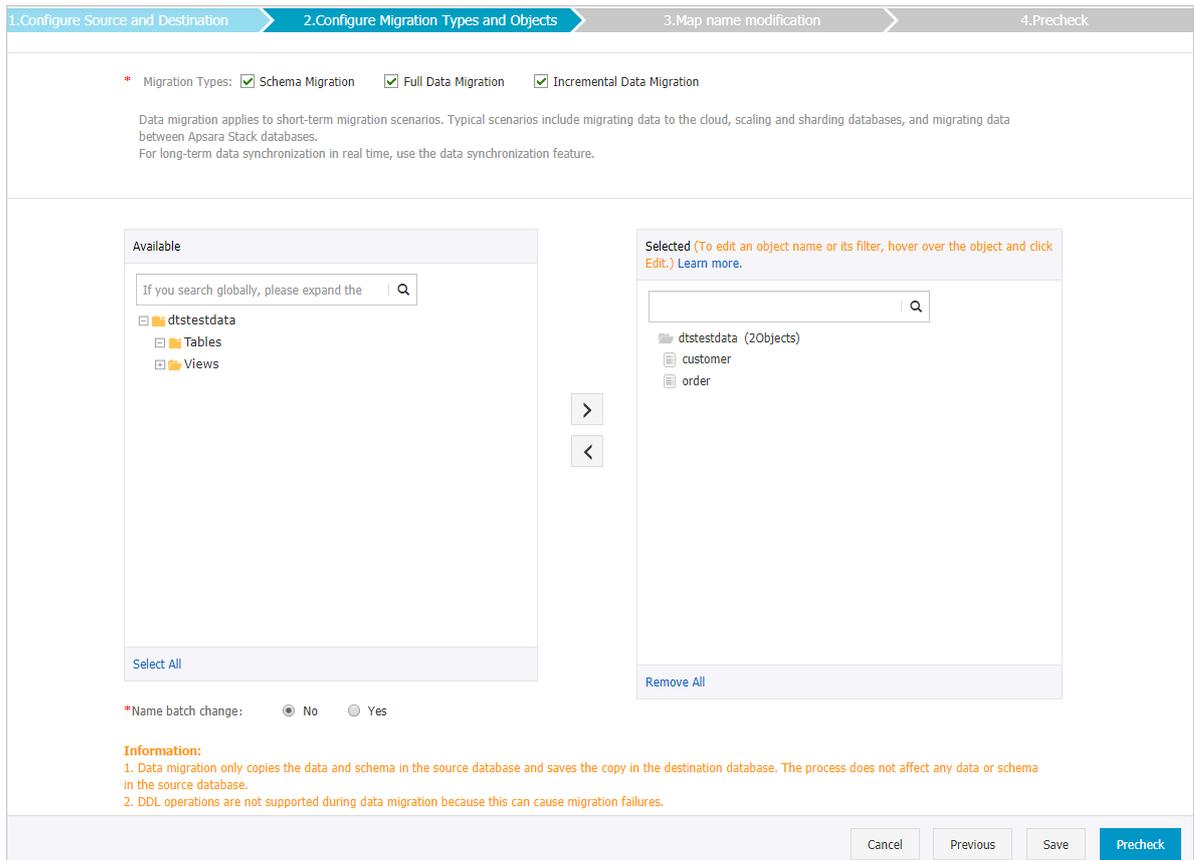
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> To perform only full data migration, select Schema Migration and Full Data Migration. If you want to migrate data without business disruptions, select Schema Migration, Full Data Migration, and Incremental Data Migration. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

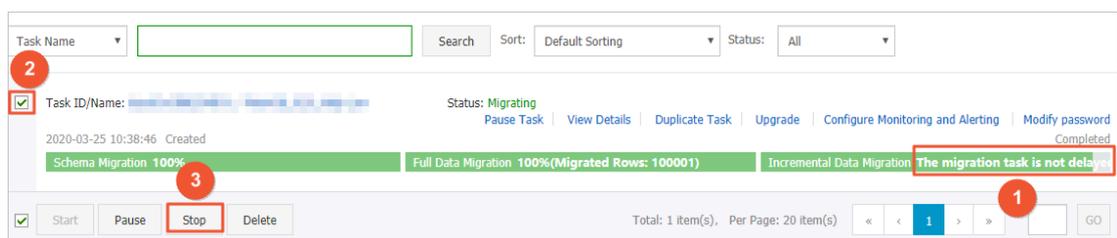
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- When 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 MySQL instance.

6.3 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.

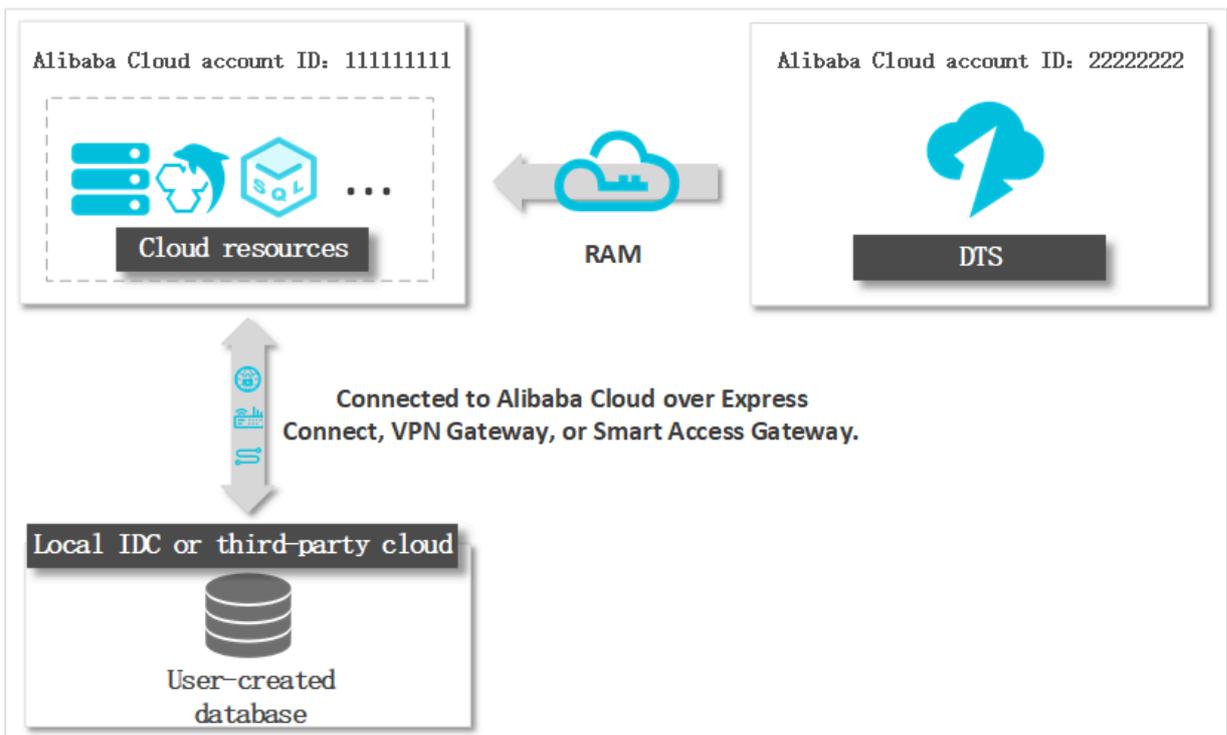
Prerequisites

- 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 [#unique_68](#).

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 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	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 #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

Migration types

- Schema migration

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



Note:

- During schema migration, DTS changes the value of the SECURITY attribute 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.

- Incremental data 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
DDL	<ul style="list-style-type: none"> 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 more information about how to create and authorize a database account, see the following topics:

- [#unique_66](#) 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

1. [#unique_66](#).
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 [#unique_71](#).

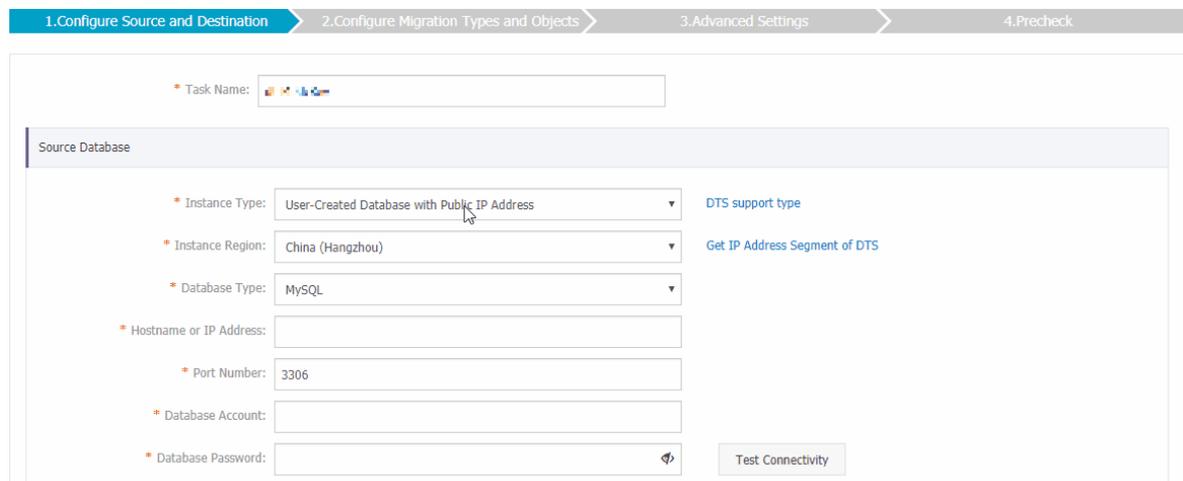
3. Create a RAM role and authorize this role to access the cloud resources of the Alibaba Cloud account to which the Express Connect circuit belongs. For more information, see [#unique_72](#).

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 **Migration Tasks** the 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.

* Task Name:

Source Database

* Instance Type: DTS support type

* Instance Region: Guide

* Apsara Stack Tenant Account ID:

* Role Name: Authorize Role Across Accounts

* Peer VPC: Proprietary network of the current login account

* Database Type:

* IP Address:

* Port Number:

* Database Account:

* Database Password: Test Connectivity

Destination Database

* Instance Type:

* Instance Region:

* RDS Instance ID:

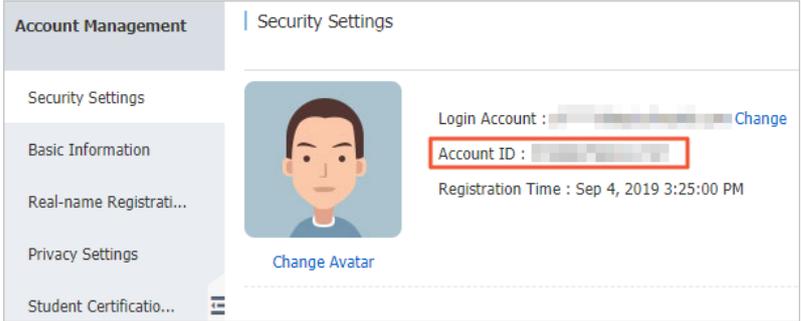
* Database Account:

* Database Password: Test Connectivity

* Encryption: Non-encrypted SSL-encrypted

Cancel **Set Whitelist and Next**

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	Select User-Created Database Connected over Express Connect, VPN Gateway, or Smart Access Gateway .
	Instance Region	Select the region of the VPC that is connected to the user-created MySQL database.

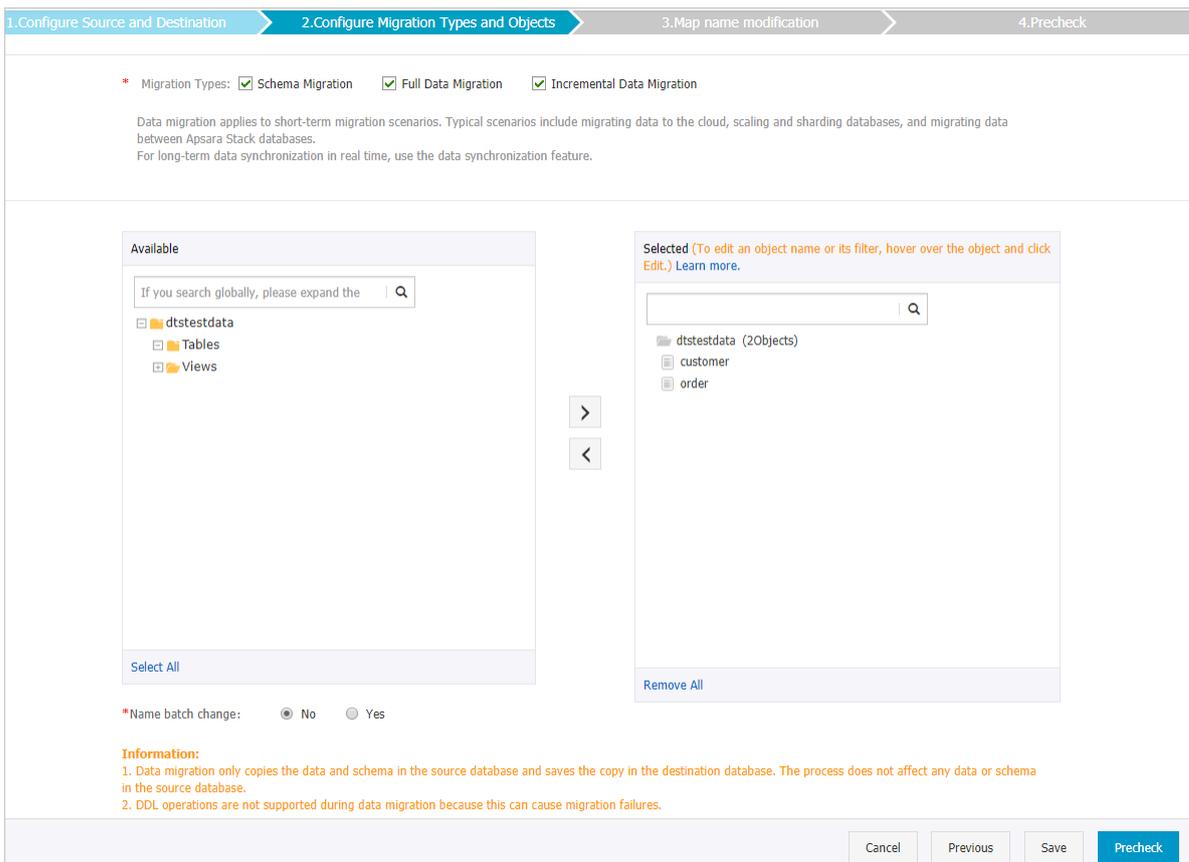
Section	Parameter	Description
	Alibaba Cloud Account ID	<p>Enter the ID of the Alibaba Cloud account to which the Express Connect circuit belongs.</p> <p> Note: To obtain the ID of the Alibaba Cloud account to which the Express Connect circuit belongs, you must log on to the Account Management console by using this account. The account ID is displayed on the Security Settings page.</p> 
	Role Name	Enter the name of the RAM role that you created earlier in Preparations .
	Peer VPC	Select the ID of the VPC that is connected to the user-created MySQL database.
	Database Type	Select MySQL .
	IP Address	Enter the endpoint that is used to access the user-created MySQL database.
	Port Number	Enter the service port number of the user-created MySQL database. The default port number is 3306 .
	Database 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 .

Section	Parameter	Description
	Database Password	<p>Enter the password of the source database account.</p> <p> 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.</p>
Destination Database	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 the permissions that are required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password of the destination database account.</p> <p> 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.</p>
	Encryption	<p>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.</p> <p> Note: The Encryption parameter is available only for regions in mainland China and the Hong Kong (China) region.</p>

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. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> 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. <div data-bbox="416 1749 1434 1944" style="background-color: #f0f0f0; padding: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

9. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

10. After the precheck is passed, click **Next**.

11. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

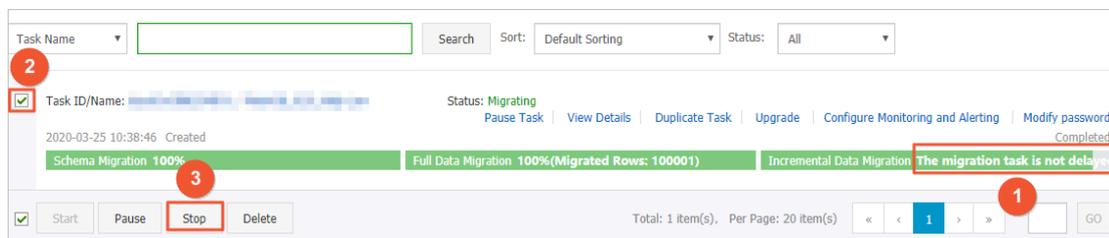
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



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

6.4 Migrate data from a user-created MySQL database to an ApsaraDB for PolarDB cluster

ApsaraDB for 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. You can use Data Transmission Service (DTS) to migrate data from a user-created MySQL database to an ApsaraDB for PolarDB cluster.

Prerequisites

- The version of the user-created MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.
- An ApsaraDB for PolarDB cluster is created. For more information, see [Create an ApsaraDB for PolarDB cluster](#).
- If the user-created MySQL database is in your local IDC, you must add the IP addresses of DTS servers to the whitelist of the database so that the servers can access your database. For more information, see [#unique_62](#).

Background information

- 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%.

- If your user-created MySQL database does not have primary key or unique constraints and each field in the database has duplicate values, the data migrated to the destination database may be duplicated.
- For columns whose data type is float or double, DTS uses the `ROUND(COLUMN, PRECISION)` function to read the values. If precision is not specified, a precision of 38 digits is set for float-type data and a precision of 308 digits is set for double-type data.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the 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 [#unique_73](#).



Note:

You can use these three migration types together to migrate data without service interruptions.

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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

SQL operations that can be synchronized during incremental data migration

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul style="list-style-type: none"> 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	SELECT privilege on objects to be migrated	SELECT, REPLICATION CLIENT, and REPLICATION SLAVE permissions on objects to be migrated.	
PolarDB for MySQL cluster	ALL permissions on migrated objects	ALL permissions on migrated objects	

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

- [#unique_66](#) for a user-created MySQL database
- [Create an account for an PolarDB for MySQL cluster](#)

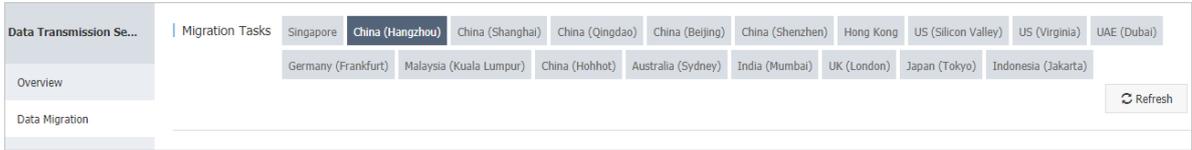
Preparations

Configure the binary logging, for information, see [#unique_66](#).

Procedure

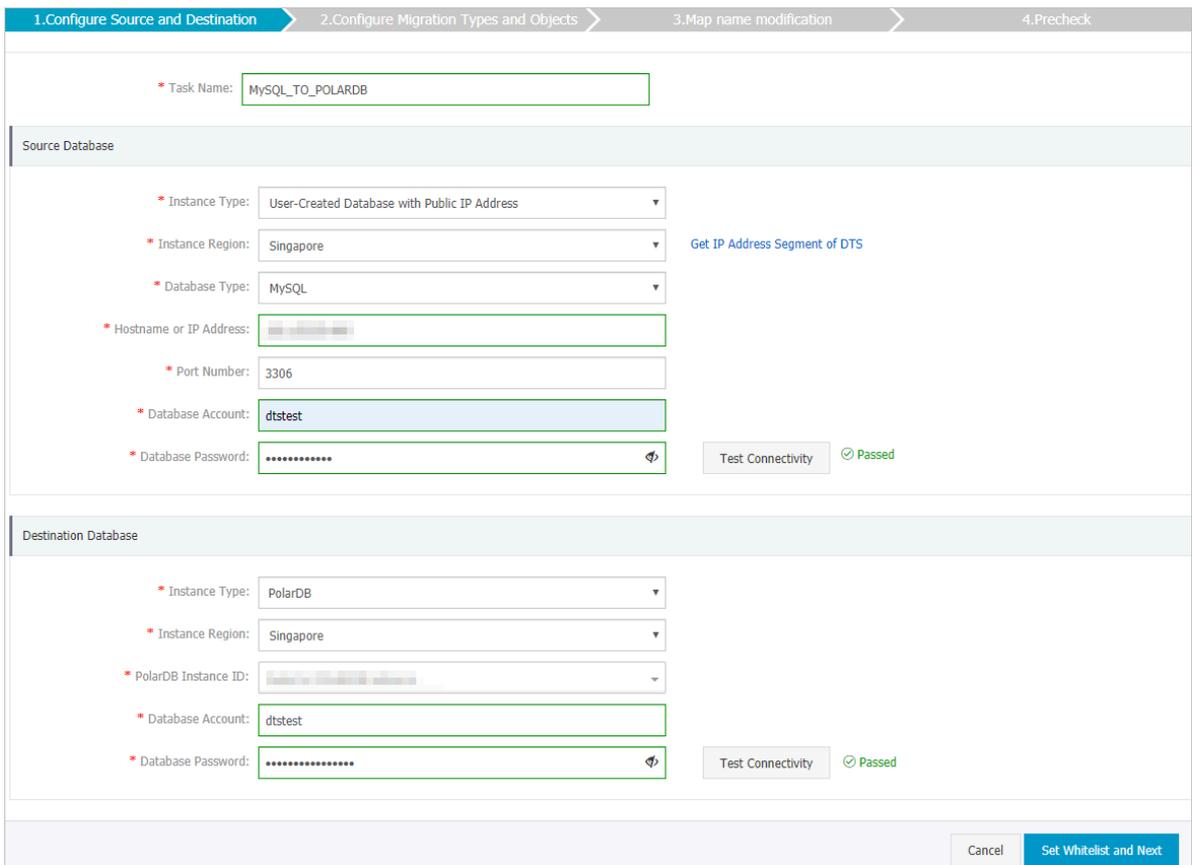
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of **Migration Tasks** the 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.

Section	Parameter	Description
Source Database	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 .  Note: If a whitelist is configured for the user-created Oracle database, you must manually add the IP addresses of DTS servers to the whitelist of the user-created Oracle database. You can click the Get IP Address Segment of DTS link next to the Instance Region parameter to obtain the IP addresses of DTS servers.
	Database Type	Select MySQL .
	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 port number configured for the user-created MySQL database. The default port number is 3306 .
	Database Account	Enter the user-created MySQL database account. For more information about permissions required for the database account, see Permissions required for database accounts .
	Database Password	Enter the password of the user-created MySQL database account.  Note: After you specify the source database parameters, click Test Connectivity next to the Database Password parameter to verify whether the parameters are valid. If the source database parameters are valid, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database parameters as prompted.
Destination Database	Instance Type	Select POLARDB .
	Instance Region	Select the region to which the ApsaraDB for PolarDB cluster belongs.

Section	Parameter	Description
	POLARDB Instance ID	Select the ID of the ApsaraDB for PolarDB cluster.
	Database Account	Enter the account of the ApsaraDB for PolarDB database. For more information about permissions required for the database account, see Permissions required for database accounts .
	Database Password	Enter the password of the ApsaraDB for PolarDB database account.  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 destination database parameters are valid, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database parameters as prompted.

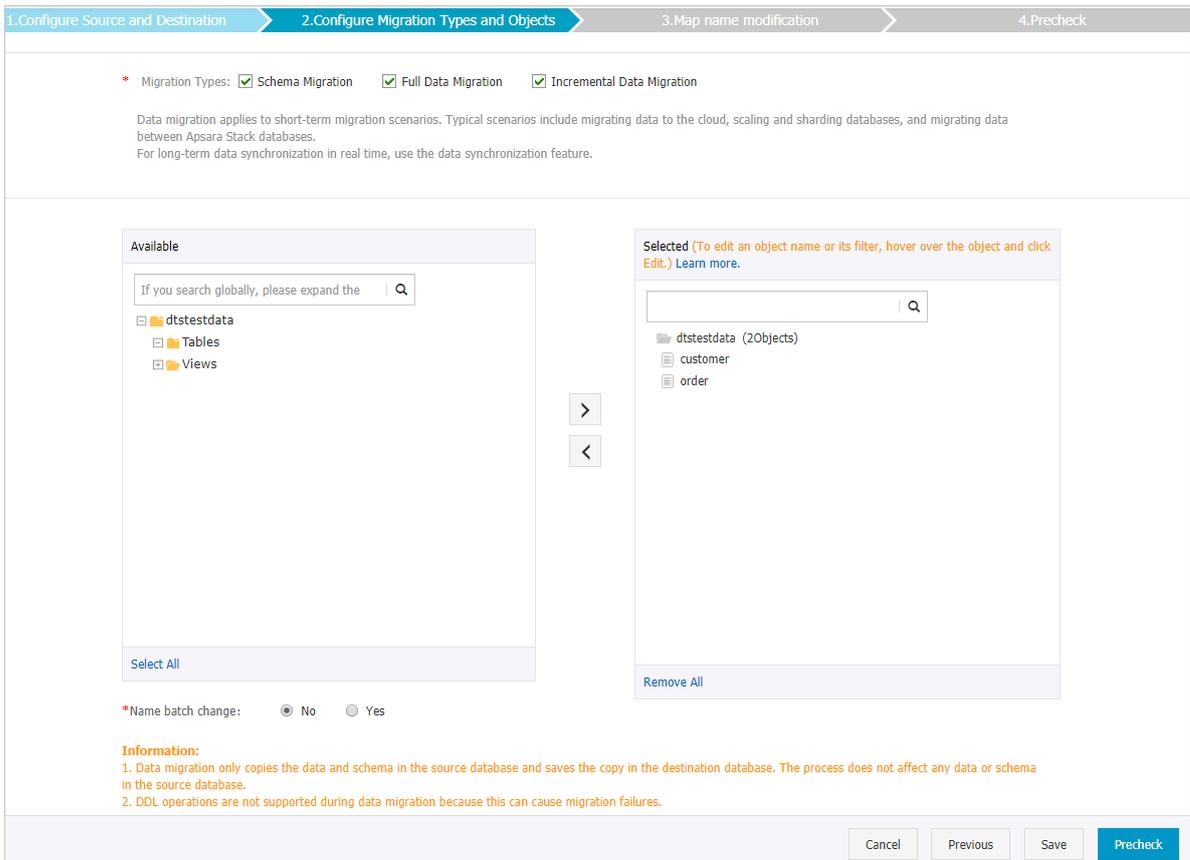
6. Click **Set Whitelist and Next** in the lower-right corner of the page.



Note:

The IP addresses of DTS servers are added to the whitelist of the ApsaraDB for PolarDB cluster. This makes sure that DTS servers can connect to the cluster.

7. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> To perform only full data migration, select Schema Migration and Full Data Migration. If you want to migrate data without business disruptions, select Schema Migration, Full Data Migration, and Incremental Data Migration. <p> 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.</p>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

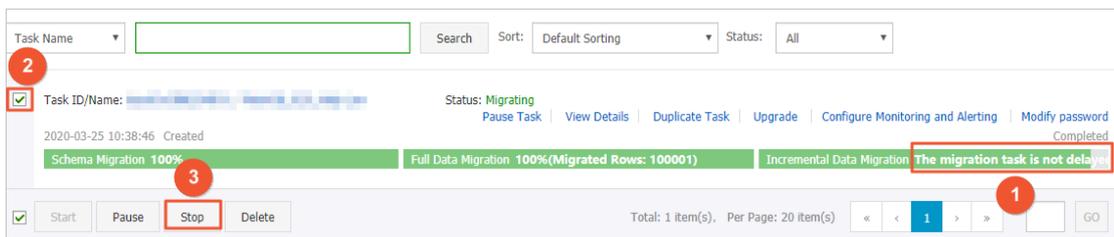
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When 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 cluster.

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

This topic describes how to use Data Transmission Service (DTS) to migrate data from a user-created MySQL database to a Distributed Relational Database System (DRDS) instance. DTS supports schema migration, full data migration, and incremental data migration.

Prerequisites

- The version of the user-created MySQL database is 5.1, 5.5, 5.6, 5.7, or 8.0.

- The database in DRDS instance are created based on ApsaraDB RDS for MySQL instances that you purchased. DTS does not support databases that are created based on private ApsaraDB RDS for MySQL instances or Apsara PolarDB for MySQL clusters.
- A DRDS instance is created based on a Relational Database System (RDS) instance. The available storage space of the RDS instance is larger than the total size of the data in the user-created MySQL database.

Background information

- DTS does not support schema migration from a user-created MySQL database to a DRDS instance.



Note:

The schema migration process migrates the schemas of the objects to be migrated 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 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.
- For columns whose data type is float or double, DTS uses the ROUND(COLUMN, PRECISION) function to read the values. If precision is not specified, a precision of 38 digits is set for float-type data and a precision of 308 digits is set for double-type data.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the task . Otherwise, the data in the source database will overwrite the data in the destination database after the task is resumed.

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 #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

Migration types

- Full data migration

DTS migrates full data of objects to be migrated from the user-created MySQL database to the destination DRDS instance.

**Note:**

Concurrent insertions are performed during full data migration. This results in table fragmentation in the destination instance. After a full data migration task is completed, the tablespace of the destination instance is larger than that of the source instance.

- Incremental data migration

During full data migration, DTS reads the binary log of the user-created MySQL database. When the full data migration is complete, DTS synchronizes the incremental data to the destination DRDS instance. This feature allows you to migrate data without interrupting your services.

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	SELECT privilege	The REPLICATION SLAVE, REPLICATION CLIENT, SHOW VIEW, and SELECT permissions
DRDS instance	Read and write permissions	Read and write permissions

Preparation

1. Configure the user-created MySQL database account and the binary log of the database. For more information, see [#unique_66](#).
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. 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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type:

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

* Database Account:

* Database Password:

Test Connectivity ✔ Passed

Destination Database

* Instance Type:

* Instance Region:

* DRDS Instance ID:

* Database Name:

* Database Account:

* Database Password:

Test Connectivity ✔ Passed

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.
Source Database	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 . <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> Note: If a whitelist is configured for the user-created SQL Server database, you must manually add the IP addresses of DTS servers to the whitelist of the user-created SQL Server database. You can click the Get IP Address Segment of DTS link next to the Instance Region parameter to obtain the IP addresses of DTS servers. Then add the IP addresses to obtain the IP addresses of DTS servers. </div>
	Database Type	Select MySQL .

Section	Parameter	Description
	Hostname or IP Address	Enter the hostname or IP address that is used to access the user-created MySQL database. In this example, enter the public IP address.
	Port Number	Enter the port number configured for the user-created MySQL database. The default port number is 3306 .
	Database Account	Enter the user-created MySQL database account. For more information about permissions required for the database account, see Permissions required for database accounts .
	Database Password	Enter the password of the user-created MySQL 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.
Destination Database	Instance Type	Select DRDS instance .
	Instance Region	Select the region to which the destination instance belongs.
	DRDS Instance ID	Select the destination DRDS instance ID.
	Database Name	Select the destination database name.
	Database Account	Enter the account of the destination database under the DRDS instance. For more information about permissions required for the database account, see Permissions required for database accounts .

Section	Parameter	Description
	Database Password	<p>Enter the password of the database account.</p> <p> 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.</p>

6. Click **Set Whitelist and Next** in the lower-right corner of the page.

**Note:**

The IP addresses of DTS servers are automatically added to the whitelist of the destination DRDS instance. Then, DTS servers can connect to the destination DRDS instance.

7. Configure migration types and objects.

* Migration Types: Full Data Migration Incremental Data Migration

Data migration applies to short-term migration scenarios. Typical scenarios include migrating data to the cloud, scaling and sharding databases, and migrating data between Apsara Stack databases. For long-term data synchronization in real time, use the data synchronization feature.

Available

If you search globally, please expand the

- data123
- mysqltest
- mysqltestnew
- sys

Select All

Selected (To edit an object name or its filter, hover over the object and click Edit.) [Learn more.](#)

- mysqltest (2Objects)
 - customer
 - vipinfo

Remove All

*Name batch change: No Yes

Information:

- Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.
- DDL operations are not supported during data migration because this can cause migration failures.

Cancel Previous Save **Precheck**

Item	Description
Migration types	<ul style="list-style-type: none"> To perform only full data migration, select Full Data Migration. If you want to migrate data without disruptions to your business, select Full Data Migration and Incremental Data Migration. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> • 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 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. </div>

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



Note:

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

- Incremental data migration

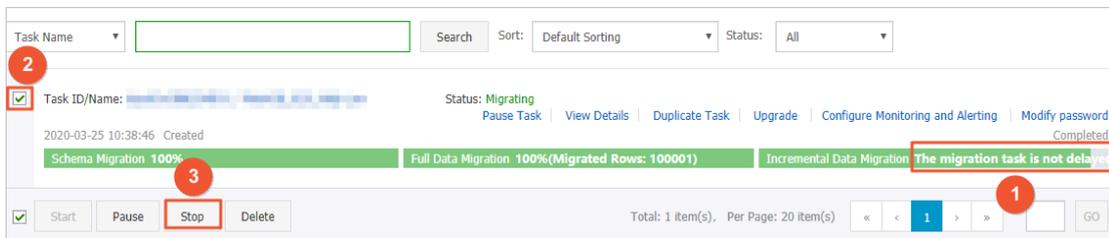
An incremental data migration task does not automatically stop. You must manually stop the migration task.



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**, manually stop the migration task.



12. Switch your workloads to the DRDS instance.

What to do next

To ensure database security, after the data migration is complete, you must delete the accounts of both the user-created MySQL database and database in the DRDS instance.

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

This topic describes how to migrate incremental data from a user-created SQL Server database to an ApsaraDB RDS for SQL Server 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 SQL Server database, you can select all of the supported migration types to ensure service continuity.

Prerequisites

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



Note:

- SQL Server Cluster and SQL Server AlwaysOn High Availability Group are not supported.

- If you migrate data between different versions of databases, make sure that the database versions are supported.

- The tables to be migrated in the user-created SQL Server database must have primary keys or UNIQUE NOT NULL indexes.
- The available storage space of the destination ApsaraDB RDS for SQL Server database is larger than the total space of the data in the user-created SQL Server database.

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%.
- To ensure the normal operation of incremental data migration, please do not frequently perform backups on the source database, and it is recommended to keep the logs for more than 3 days to avoid the inability to obtain logs after truncation.
- To ensure the accuracy of incremental data migration, DTS adds a heartbeat table to the user-created SQL Server database. The table name is `_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 manually 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](#).

- 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.

Limits

- DTS does not support the schema migration of assemblies, service brokers, full-text indexes, full-text catalogs, distributed schemas, distributed functions, CLR stored procedures, CLR scalar functions, CLR value functions, internal tables, systems, and aggregate functions.
- Data of the `sql_variant` type cannot be migrated.
- Tables that contain computed columns cannot be migrated.
- 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 database.

- Incremental data migration

After full data migration is complete, DTS synchronizes incremental data from the user-created SQL Server database to the destination ApsaraDB RDS for SQL Server database. Incremental data migration helps you 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 operations

**Note:**

The UPDATE operations that update only the large fields cannot be synchronized.

- CREATE TABLE operations

**Note:**

The CREATE TABLE operations for creating partition tables or tables that contain functions cannot be synchronized.

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

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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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 database	The read/write permissions	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 a user-created SQL Server database
- [Create an account for an RDS for SQL Server instance](#) for an ApsaraDB RDS for SQL Server database

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 incremental data migration.

**Note:**

Before performing incremental data migration, do not perform DDL operations on the objects to be migrated in the user-created SQL Server database. Otherwise, the objects may fail to be migrated.

Preparations before data migration

Before configuring a data migration task, you must configure logs 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 mode:

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

Parameters:

<database_name>: the name of the source database.

Example:

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

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

```
BACKUP DATABASE <database_name> TO DISK='<physical_backup_device_name>';  
GO
```

Parameters:

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

Example:

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

```
GO
```

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

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

Parameters:

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

Example:

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

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

1. Configure Source and Destination
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type:

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

* Database Account:

* Database Password: ✔ Passed

Destination Database

* Instance Type:

* Instance Region:

* RDS Instance ID:

* Database Account:

* Database Password: ✔ Passed

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 Database	Instance Type	<p>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.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.</p> <p> Note: If a whitelist is configured for the user-created SQL Server database, you must manually add the CIDR blocks of DTS servers to the whitelist of the user-created SQL Server database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.</p>
	Database 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 .
	Database Account	Enter the account for the user-created SQL Server database. For more information about permissions required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password for the database account.</p> <p> 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.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination RDS instance resides.
	RDS Instance ID	Select the ID of the destination RDS instance.

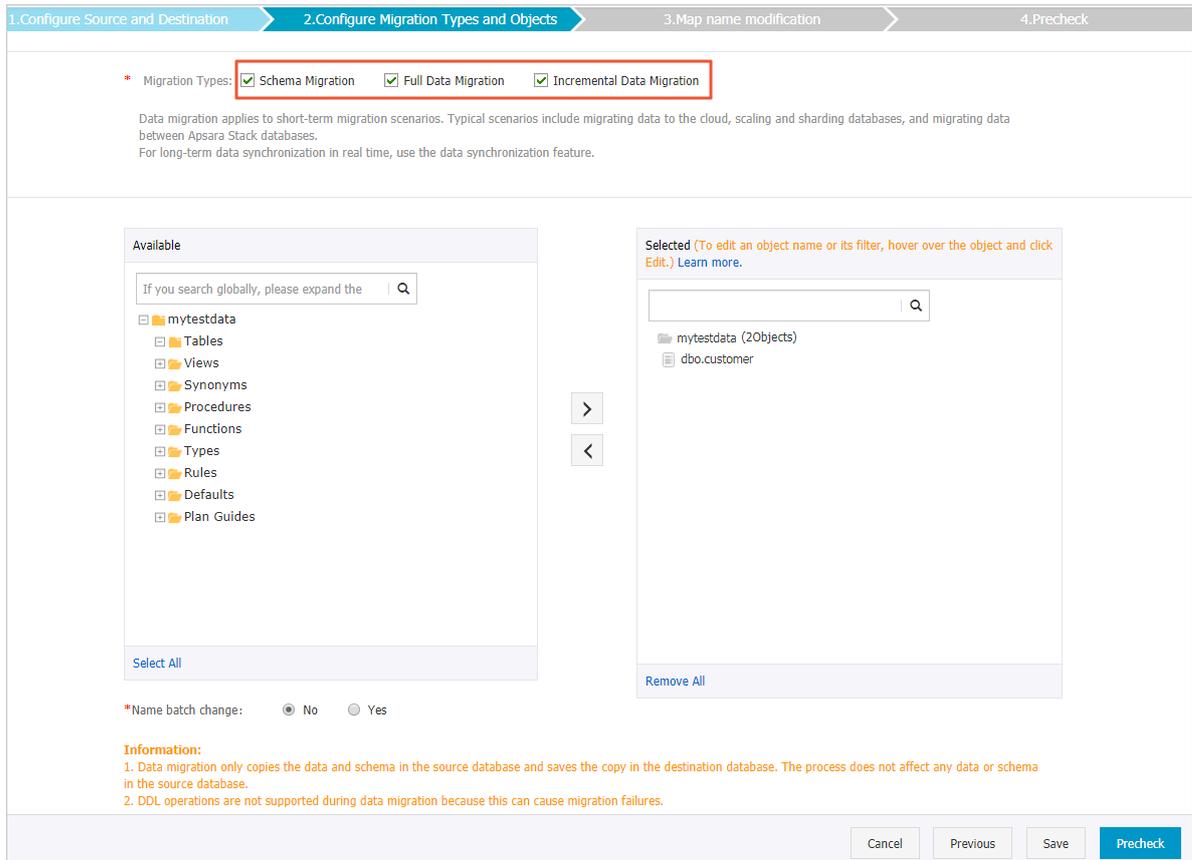
Section	Parameter	Description
	Database Account	Enter the database account of the destination RDS instance. For more information about permissions required for the account, see Permissions required for database accounts .
	Database Password	Enter the password for the database account.  Note: After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the 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.



Parameter	Description
Migration types	<ul style="list-style-type: none"> 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. <p>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.</p> <p>In this example, you must select Schema Migration, Full Data Migration, and Incremental Data Migration.</p>

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> • 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 databases, tables, or columns 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 SQL Server database. If you want an object to have a different name after the object is migrated to the destination RDS 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**.



Note:

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

- Incremental data migration

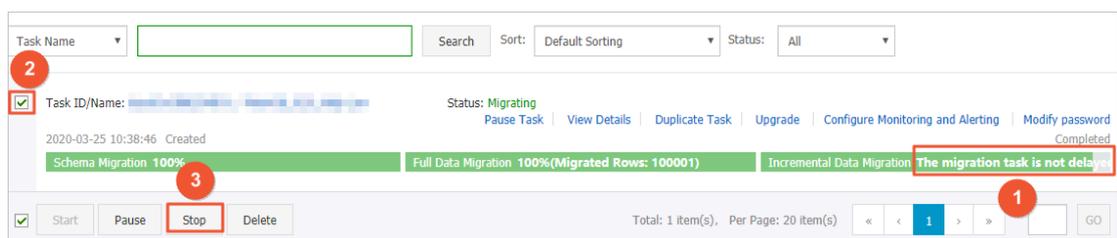
An incremental data migration task does not automatically stop. You must manually stop the migration task.



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.

- 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.
- 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 SQL Server database and ApsaraDB RDS for SQL Server database to ensure database security.

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

This topic describes how to migrate full data from a user-created SQL Server database to an ApsaraDB RDS for SQL Server instance by using Data Transmission Service (DTS). DTS supports schema migration, full data migration, and incremental data migration. To migrate full data from a user-created SQL Server database, you can select schema migration and full data migration.

**Note:**

For more information about how to migrate data without service disruptions, see [Migrate incremental data from a user-created SQL Server database to an ApsaraDB RDS for SQL Server database](#).

Prerequisites

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

**Note:**

- 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 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 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 manually 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](#).

- 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 #unique_51 .

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:

- [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 database

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

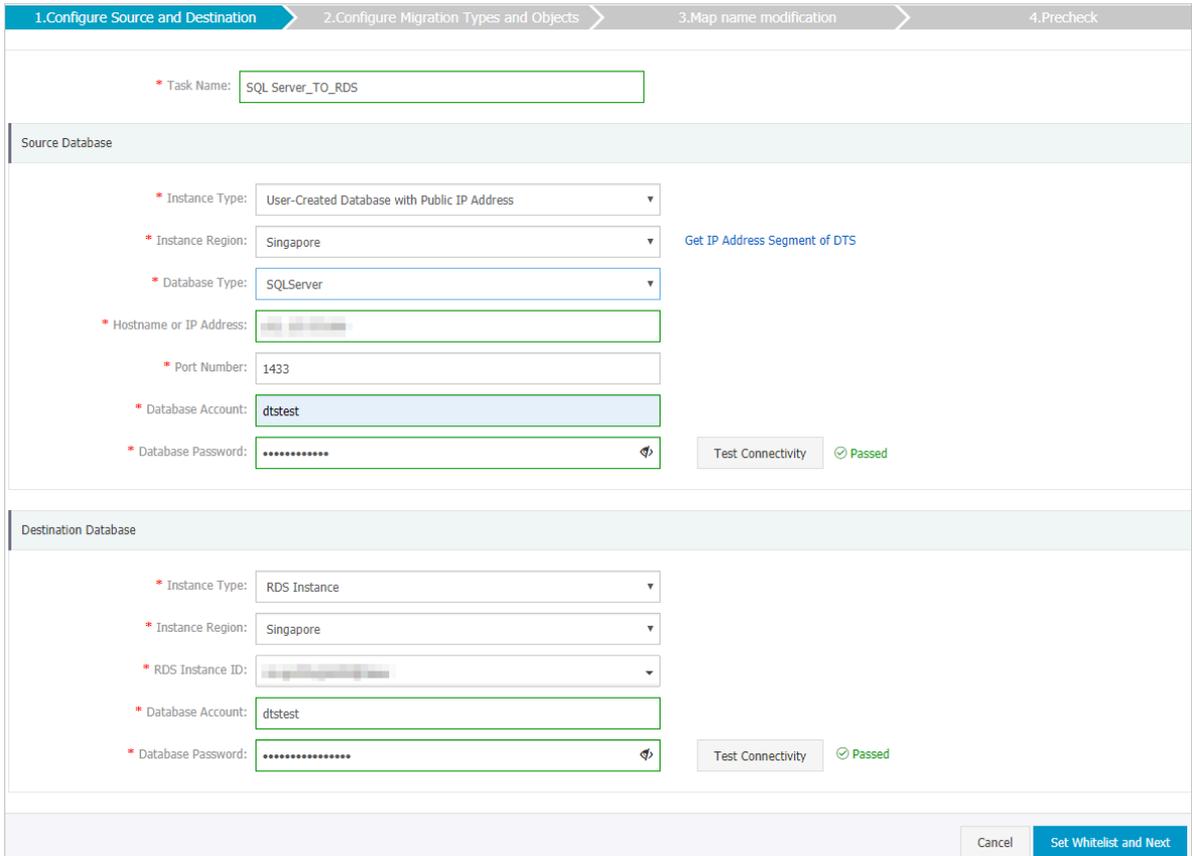
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. 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.



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.
Source Database	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. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If Instance Type is set to User-Created Database with Public IP Address, you do not need to specify the Instance Region.</p> <p> 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.</p>
	Database 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	<p>Enter the service port number of the user-created SQL Server database. The default port number is 1433.</p> <p> Note: In this example, the service port of the user-created MongoDB database must be open to the public network.</p>
	Database Account	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 .
Destination Database	Database Password	<p>Enter the password of the source database account.</p> <p> 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.</p>
	Instance Type	Select RDS Instance .
	Instance Region	Select the region in which the destination RDS instance resides.

Section	Parameter	Description
	RDS Instance ID	Select the ID of the destination RDS instance.
	Database 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 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 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.

The screenshot shows the '2. Configure Migration Types and Objects' step of a migration process. At the top, there are four tabs: '1. Configure Source and Destination', '2. Configure Migration Types and Objects' (active), '3. Map name modification', and '4. Precheck'. Below the tabs, the 'Migration Types' section has three options: 'Schema Migration' (checked), 'Full Data Migration' (checked), and 'Incremental Data Migration' (unchecked). A note below states: 'During full data migration, data updates in the source database are not migrated to the destination instance. For data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration.' The main area is split into two panels. The 'Available' panel on the left shows a tree view of database objects under 'mytestdata', including Tables, Views, Synonyms, Procedures, Functions, Types, Rules, Defaults, and Plan Guides. The 'Selected' panel on the right shows 'mytestdata (20Objects)' and 'dbo.customer'. Navigation arrows are between the panels. At the bottom, there are radio buttons for 'Name batch change' (No selected, Yes unselected) and an 'Information' section with two points: '1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.' and '2. DDL operations are not supported during data migration because this can cause migration failures.' At the very bottom are buttons for 'Cancel', 'Previous', 'Save', and 'Precheck'.

Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <p>In this example, select Schema Migration and Full Data Migration.</p> <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p> Note: To ensure data consistency, we recommend that you do not write new data into the user-created SQL Server database during full data migration.</p> </div>

Parameter	Description
Migration objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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. In the lower-right corner of the page, click **Precheck**.



Note:

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

6.8 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 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 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

Migration type	Description
Schema migration	<p>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.</p> <p> Note: However, if an object contains triggers, the data between the source and destination databases will become inconsistent.</p>
Full data migration	<p>DTS migrates historical data of the required objects from the user-created Oracle database to the destination PolarDB cluster.</p> <p> 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.</p>
Incremental data migration	<p>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.</p>

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.

Procedure

1. Log on to the [DTS console](#).

- In the left-side navigation pane, click **Data Migration**.
- At the top of **Migration Tasks** the page, select the region where the destination cluster resides.



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

* Task Name:

Source Database

* Instance Type: [DTS support type](#)

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

* Instance Type: Non-RAC Instance RAC or PDB Instance

* SID:

* Database Account:

* Database Password:

Destination Database

* Instance Type:

* Instance Region:

* PolarDB Instance ID:

* Database Name:

* Database Account:

* Database Password:

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
Source Database	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 #unique_67 .
	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 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.
	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	<ul style="list-style-type: none"> • 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.
	Database Account	Enter the account of the user-created Oracle database. For more information about the permissions that are required for the account, see Preparations .

Section	Parameter	Description
	Database Password	<p>Enter the password for the source database account.</p> <p> 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.</p>
Destination 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.
	Database Name	Enter the name of the destination database.
	Database Account	Enter the database account of the destination PolarDB cluster. For more information about the permissions that are required for the account, see Preparations .
	Database Password	<p>Enter the password for the destination database account.</p> <p> 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.</p>

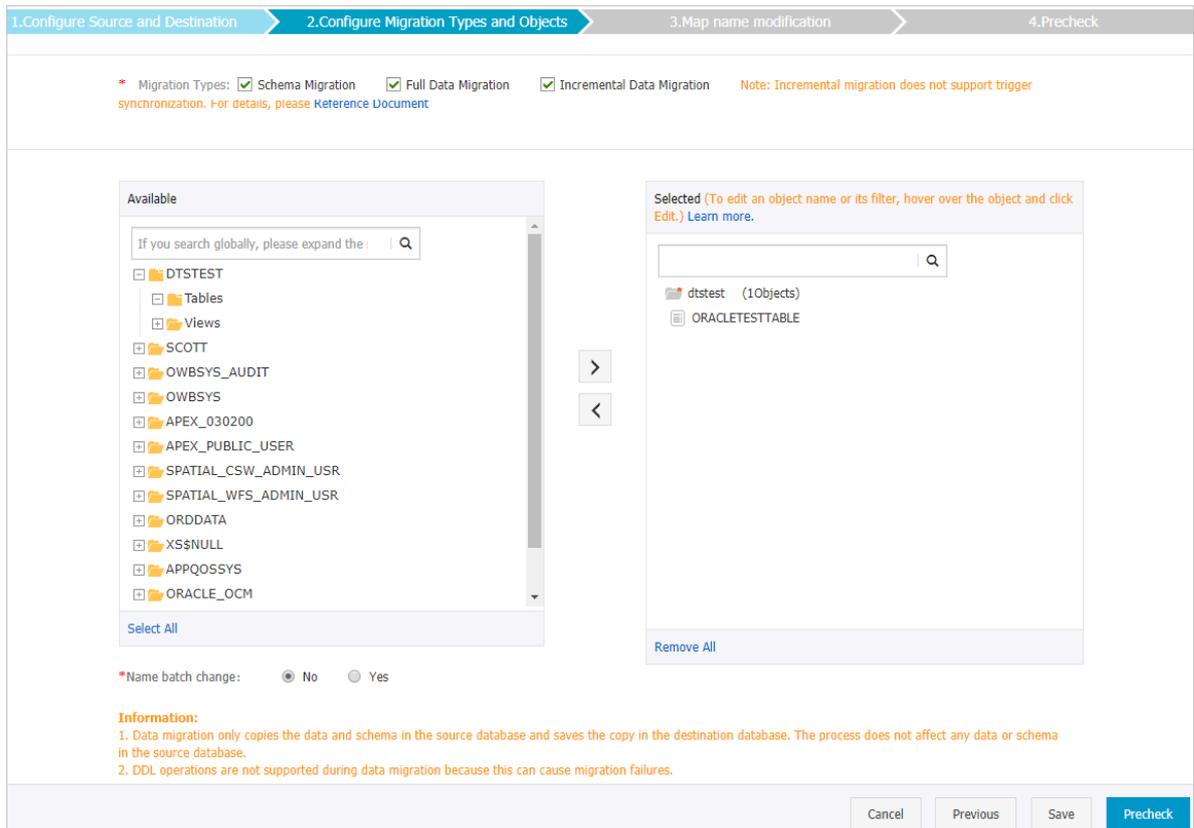
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <p> 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.</p>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

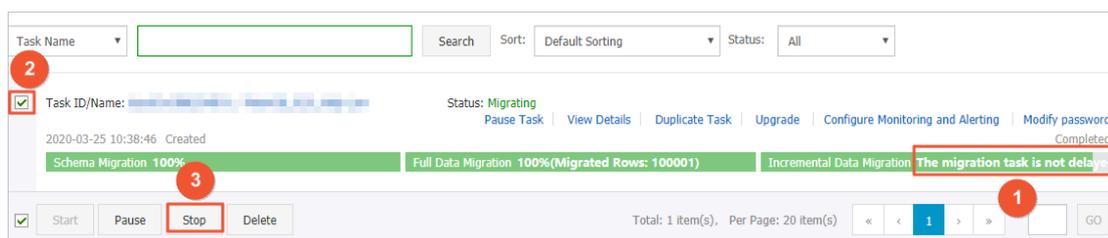
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- When 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 cluster.

6.9 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 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.
- 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 #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

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 data migration

DTS migrates historical data of the required objects from the user-created Oracle database to the destination ApsaraDB RDS for MySQL instance.

- 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 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:

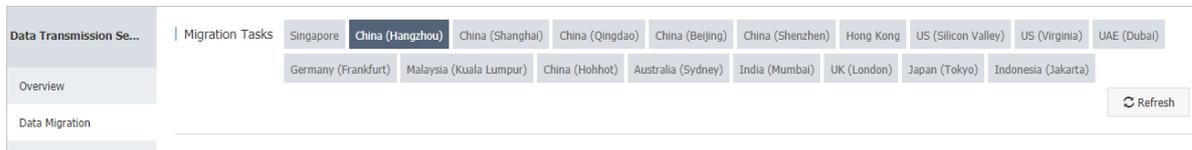
- [CREATE USER](#) and [GRANT](#) for a user-created Oracle database
- [Create an account for an RDS for MySQL instance](#) and [Change the permissions of an account for an RDS for MySQL instance](#)

**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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type: DTS support type

* Instance Region: Get IP Address Segment of DTS

* Database Type:

* Hostname or IP Address:

* Port Number:

* Instance Type: Non-RAC Instance RAC Instance

* SID:

* Database Account:

* Database Password:

Destination Database

* Instance Type:

* Instance Region:

* RDS Instance ID:

* Database Account:

* Database Password:

* Encryption: Non-encrypted SSL-encrypted

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.
Source Database	Instance Type	<p>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.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.</p> <p> 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.</p>
	Database Type	Select Oracle .
	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	<ul style="list-style-type: none"> • 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.
	Database Account	Enter the account of the user-created Oracle database. For more information about the permissions that are required for the account, see Preparations .
	Database Password	<p>Enter the password for the source database account.</p> <p> 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.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination RDS instance resides.

Section	Parameter	Description
	RDS Instance ID	Select the ID of the destination RDS instance.
	Database Account	Enter the database account of the destination RDS instance. For more information about the permissions that are required for the account, see Preparations .
	Database Password	Enter the password for the destination database account.  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.

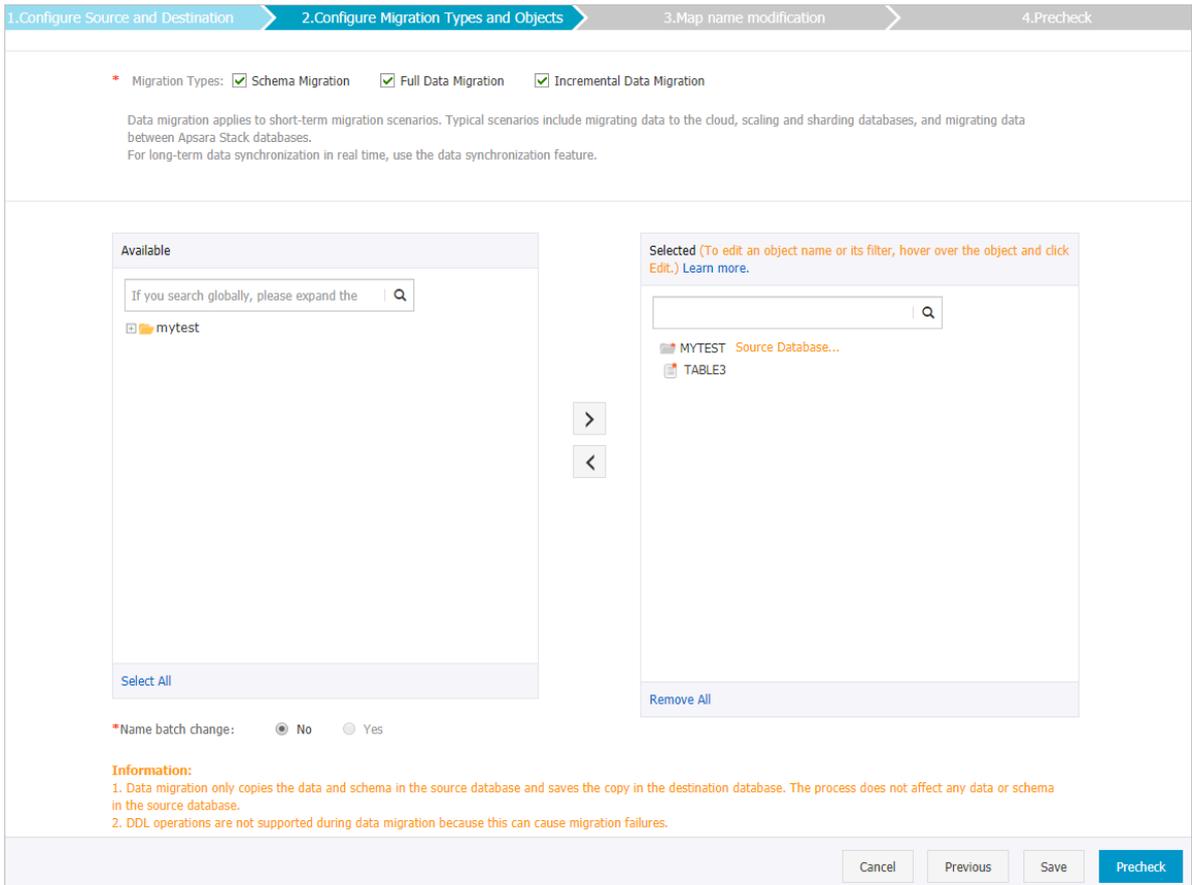
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <p>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.</p>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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**.



Note:

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

- Incremental data migration

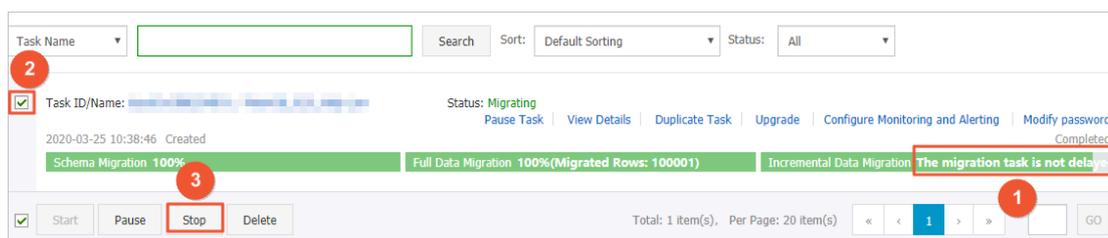
An incremental data migration task does not automatically stop. You must manually stop the migration task.



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.

- 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.
- 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

e data changes from the ApsaraDB RDS for MySQL instance to the user-created Oracle database. To do this, submit a ticket.

6.10 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 that you purchased. DTS does not support databases that are created based on private ApsaraDB RDS for MySQL instances or 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 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 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

- Full data 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.

- Incremental data 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

1. Create a database and tables in the destination DRDS instance based on the schema of the user-created 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](#).



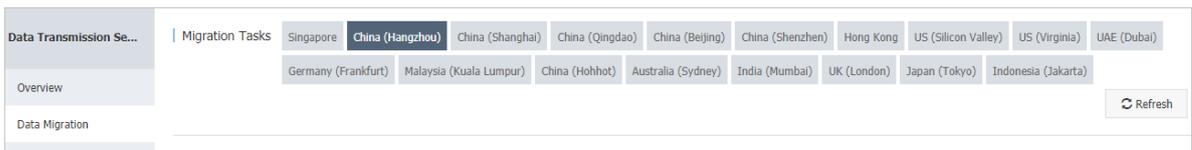
Notice:

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. 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
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
Source Database	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 #unique_67 .
	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 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 .
	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	<ul style="list-style-type: none"> • 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.
	Database Account	Enter the account of the user-created Oracle database. For more information about the permissions that are required for the account, see Preparations .

Section	Parameter	Description
	Database Password	<p>Enter the password for the source database account.</p> <p> 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.</p>
Destination Database	Instance Type	Select DRDS Instance .
	Instance Region	Select the region where the destination DRDS instance resides.
	DRDS Instance ID	Select the ID of the destination DRDS instance.
	Database 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 Password	<p>Enter the password for the destination database account.</p> <p> 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.</p>

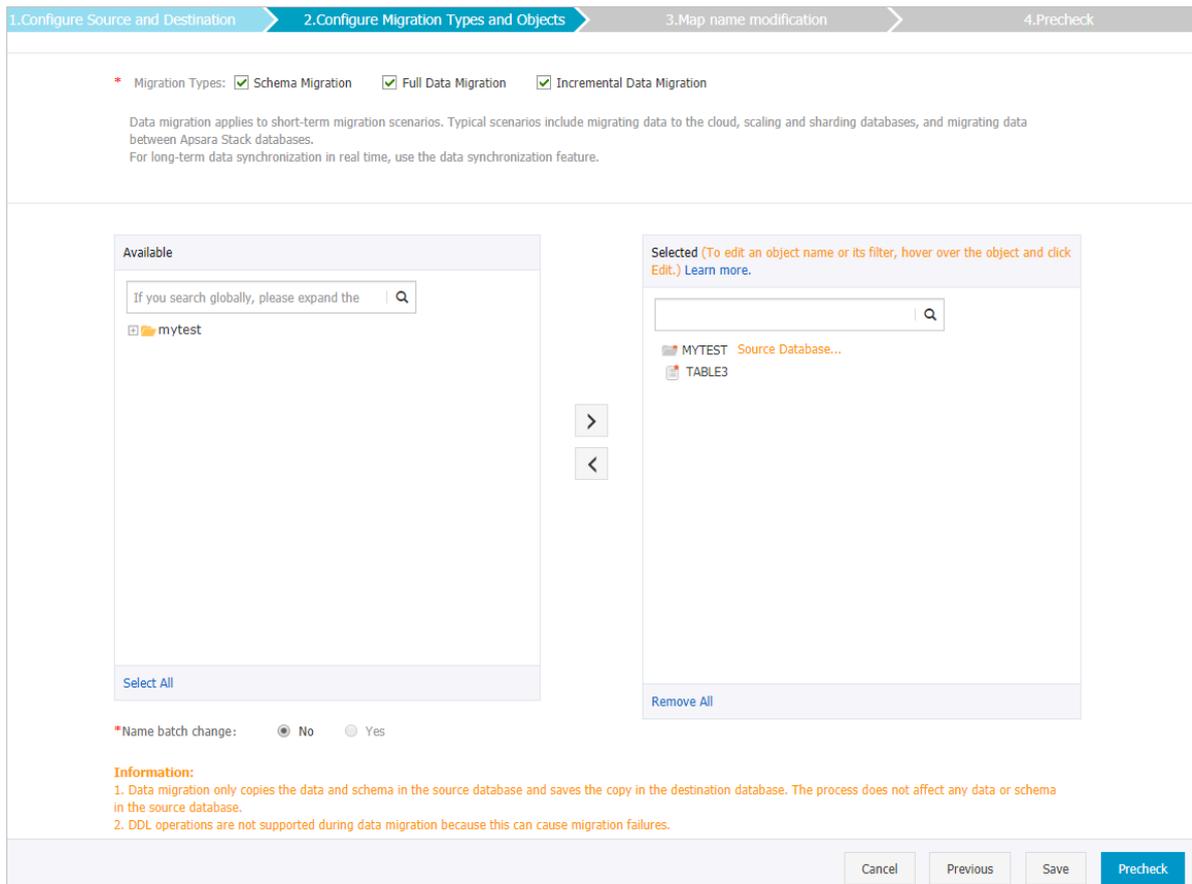
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



Note:

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

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



Parameter	Description
Migration Types	<ul style="list-style-type: none"> To perform only full data migration, select Full Data Migration. To migrate data with minimal downtime, select Full Data Migration and Incremental Data Migration. <p> 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.</p>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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.

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



Note:

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

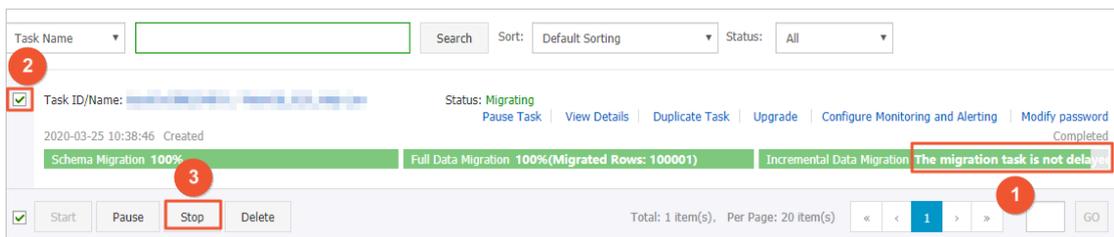
Do not manually stop a task during full data migration. Otherwise, data migrated to the destination database will be incomplete. Wait until the migration task automatically stops.

- 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**, 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.

6.11 Migrate data from a user-created Oracle database to an ApsaraDB RDS for PPAS instance

This topic describes how to migrate data from a user-created Oracle database to an ApsaraDB RDS for PPAS 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 to an ApsaraDB RDS for PPAS instance, 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.



Note:

We recommend that you use ApsaraDB RDS for PPAS V10 for higher compatibility.

- 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 available storage space of the destination ApsaraDB RDS for PPAS instance is larger than the total size of the data in the user-created Oracle database.
- The service port of the user-created Oracle 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 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 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

- During schema migration, the reverse indexes and bitmap indexes of the source database are stored as common indexes in the ApsaraDB RDS for PPAS instance.
- During schema migration, partitioned indexes are converted into independent indexes on each partition in the ApsaraDB RDS for PPAS instance.
- Incremental data migration supports only tables that have primary keys or UNIQUE NOT NULL indexes.
- Incremental data migration does not support the LONG data type.
- DDL operations that are performed during incremental data migration cannot be synchronized to the destination database.
- Materialized views cannot be migrated.

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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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, synonym, trigger, stored procedure, function, package, and user-defined type.

- Full data migration

DTS migrates historical data of the required objects from the user-created Oracle database to the destination ApsaraDB RDS for PPAS instance.

**Note:**

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

- Incremental data 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 destination ApsaraDB RDS for PPAS instance. Incremental data migration allows you to ensure service continuity when you migrate data from the user-created Oracle database to the destination ApsaraDB RDS for PPAS instance.

Data type mappings

For more information, see [Data type mappings between heterogeneous databases](#).

Preparations

**Note:**

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

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

Database	Schema migration	Full data migration	Incremental data migration
User-created Oracle database	The owner permission on schemas	The owner permission on schemas	The database administrator (DBA) permission
ApsaraDB RDS for PPAS instance	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 the user-created Oracle database

- [Create an account](#) for the ApsaraDB RDS for PPAS 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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type:

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

* Instance Type: Non-RAC Instance RAC Instance

* SID:

* Database Account:

* Database Password:

Destination Database

* Instance Type:

* Instance Region:

* RDS Instance ID:

* Database Name:

* Database Account:

* Database Password:

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.
Source Database	Instance Type	<p>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.</p> <div style="background-color: #f2f2f2; padding: 10px; border: 1px solid #ccc;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.</p> <p> 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.</p>
	Database Type	Select Oracle .
	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	<ul style="list-style-type: none"> • 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.
	Database Account	Enter the account of the user-created Oracle database. For more information about the permissions that are required for the account, see Preparations .
	Database Password	<p>Enter the password for the source database account.</p> <p> 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.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination ApsaraDB RDS for PPAS instance resides.

Section	Parameter	Description
	RDS Instance ID	Select the ID of the destination ApsaraDB RDS for PPAS instance.
	Database Name	Enter the name of the destination database.
	Database Account	Enter the database account of the destination ApsaraDB RDS for PPAS instance. For more information about the permissions that are required for the account, see Preparations .
	Database Password	Enter the password for the destination database account.  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.

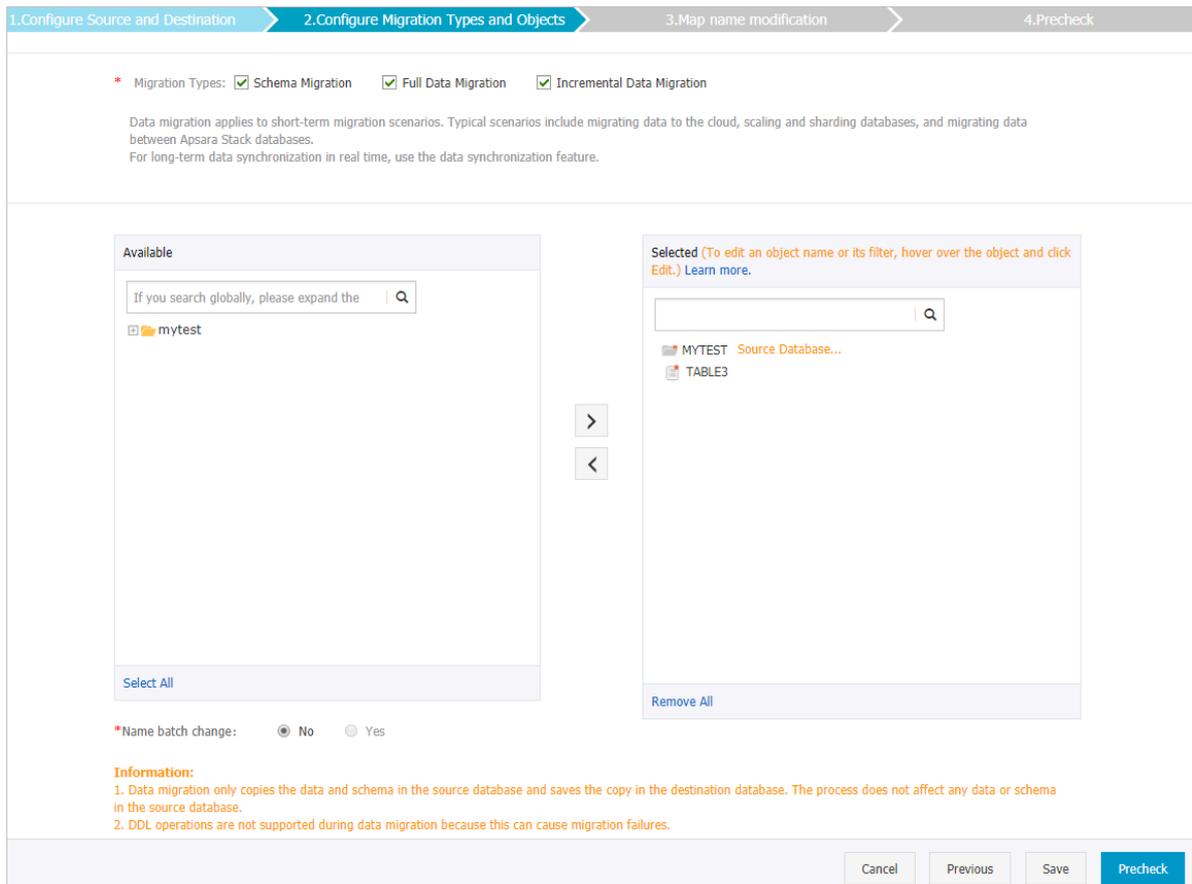
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



Note:

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

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



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> 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. Incremental data migration supports only tables that have primary keys or UNIQUE NOT NULL indexes. Incremental data migration does not support the LONG data type. </div>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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.

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



Note:

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

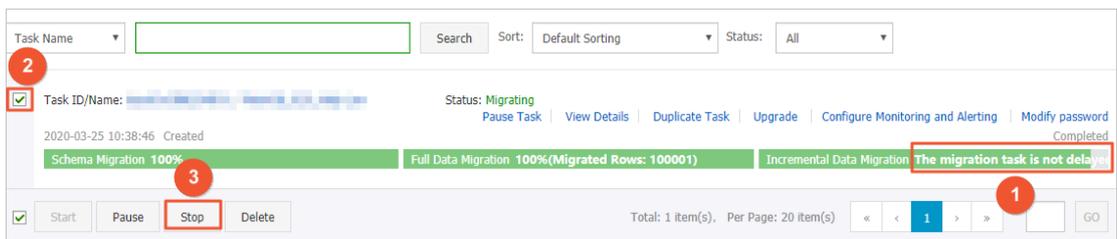
Do not manually stop a task during full data migration. Otherwise, data migrated to the destination database will be incomplete. Wait until the migration task automatically stops.

- 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**, manually stop the migration task.



12. Switch your workloads to the destination ApsaraDB RDS for PPAS 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 PPAS instance to ensure database security.

6.12 Migrate incremental data from a user-created PostgreSQL database (version 10.x 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.x, 11, or 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 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.
- To ensure that the delay time of incremental data migration is accurate, DTS adds a heartbeat table named `dts_postgres_heartbeat` 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 on to the secondary database to manually clear the replication slot.

Query Editor Query History
Scratch Pad ✕

```
1 SELECT * FROM pg_replication_slots;
```

Data Output
Explain
Messages
Notifications

slot_name	plugin name	slot_type	datoid	database name	temporary	active	active_pid
name	name	text	oid	name	boolean	boolean	integer
1 dts_sync_ohu	pgoutput	logical	16	dtstestdata	false	true	

- 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

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

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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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

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 [#unique_73](#).

Data migration process	Description
1. Schema migration	DTS migrates the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates to the destination database. <div data-bbox="539 1704 608 1774" data-label="Image"> </div> Note: Functions that are written in the C programming language cannot be migrated.
2. Full data migration	DTS migrates historical data of the required objects to the destination database.
3. Schema migration	DTS migrates the schemas of triggers and foreign keys to the destination database.

Data migration process	Description
4. Incremental data migration	DTS migrates incremental data of the required objects to the destination database. <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: Incremental data migration does not support the bit data type. </div>

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`.

```
# - Settings -
wal_level = logical                                # minimal, replica, or logical
                                                    # (change requires restart)
```



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 [#unique_62](#).



Note:

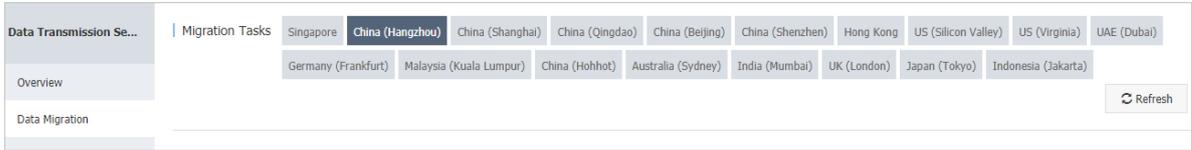
For more information about how to configure the `pg_hba.conf` file, see [The pg_hba.conf File](#). Skip this step if you have set the IP address in the `pg_hba.conf` file to `0.0.0.0/0`.

```
# "local" is for Unix domain socket connections only
local all all peer
# IPv4 local connections:
host all all 0.0.0.0/0 trust
# IPv6 local connections:
host all all ::1/128 md5
```

Procedure

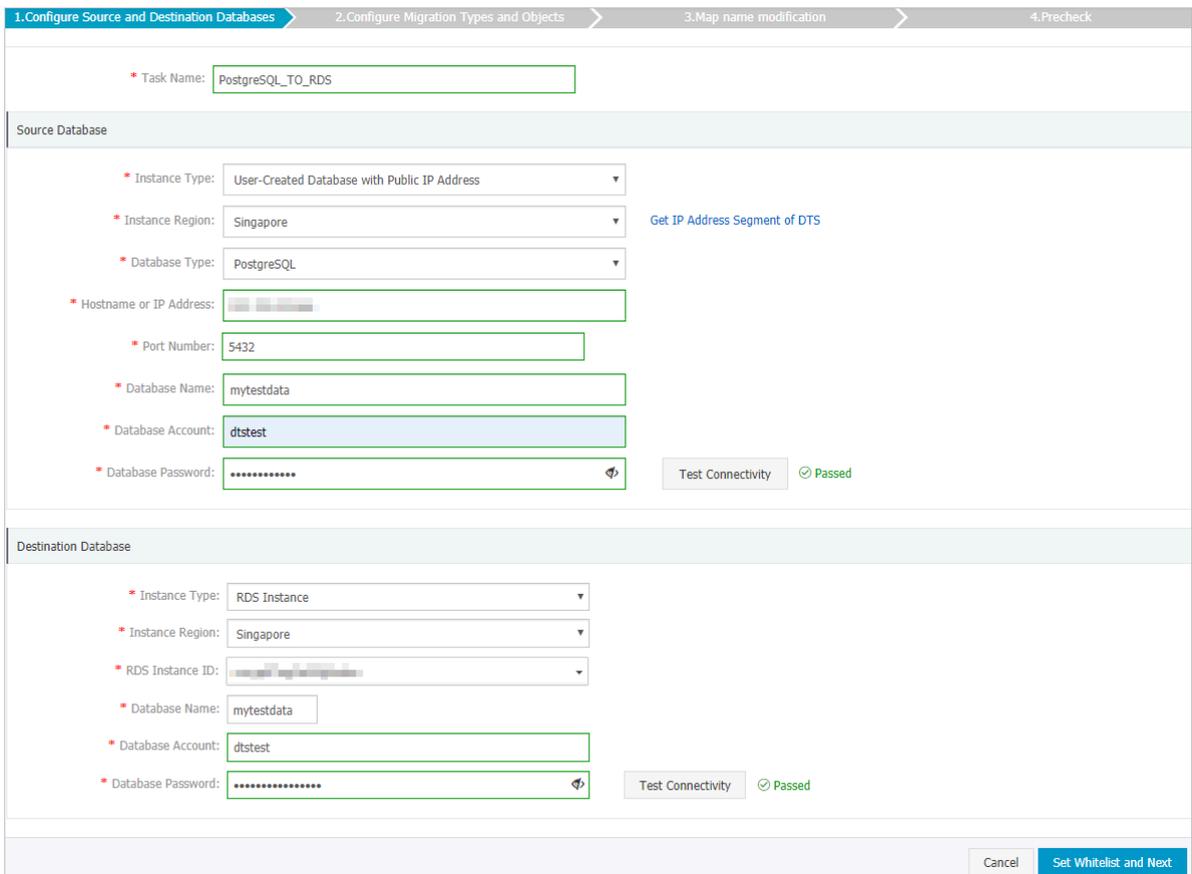
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of **Migration Tasks** the 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.

Section	Parameter	Description
Source Database	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 #unique_67 .
	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.
	Database Name	Enter the name of the user-created PostgreSQL database.
	Database Account	Enter the account of the user-created PostgreSQL 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 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
Destination Database	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 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 in the ApsaraDB RDS for PostgreSQL instance. For more information, see Create a database .
	Database 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 .
	Database Password	Enter the password for the destination database account.  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.

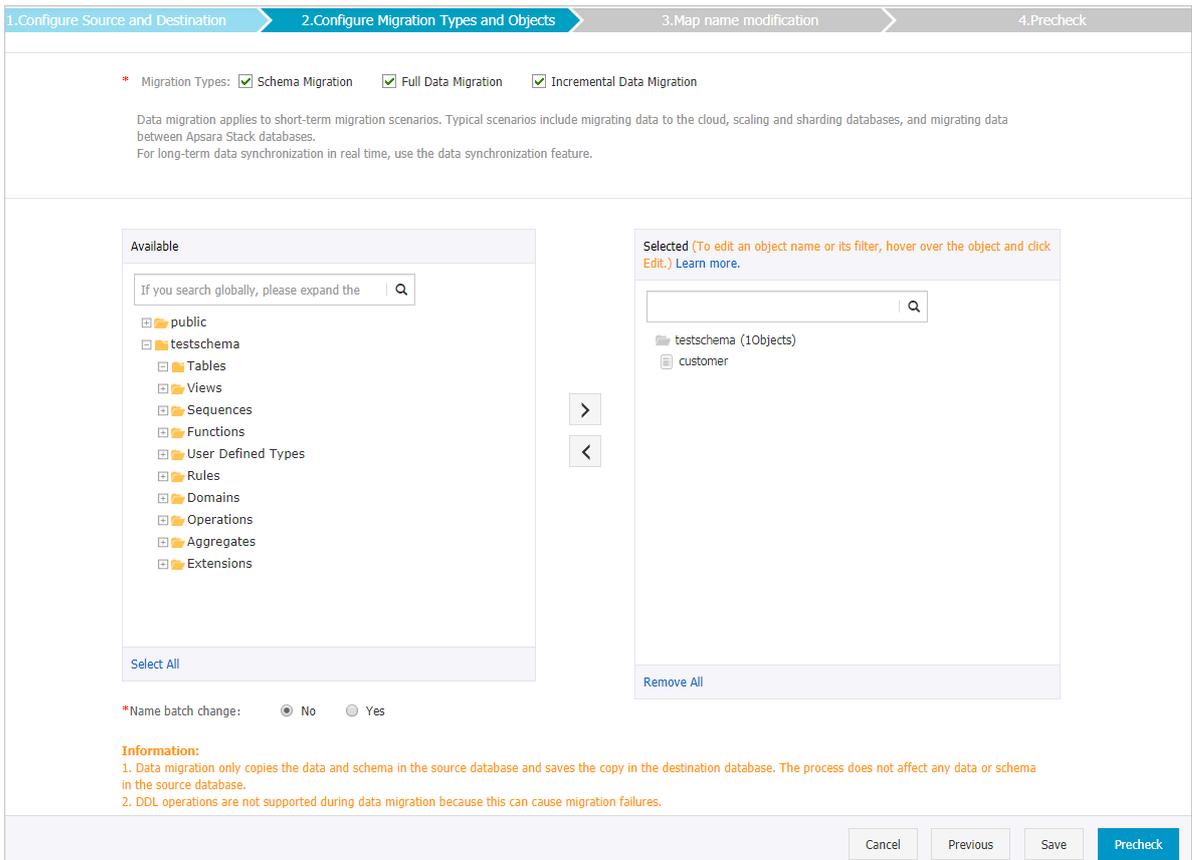
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



Note:

The CIDR blocks of DTS servers are automatically added to the whitelist of the destination 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.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p> 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.</p> </div>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> You can select columns, tables, or databases as the objects to be migrated. If you select a database 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: <pre>ALTER TABLE schema.table REPLICA IDENTITY FULL;</pre> <p>You must replace the schema and table in the preceding sample statement with the actual schema name and table name.</p> 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. 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. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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 [#unique_59](#). If you do not need to switch your workloads, you can perform the following steps to stop the migration task.

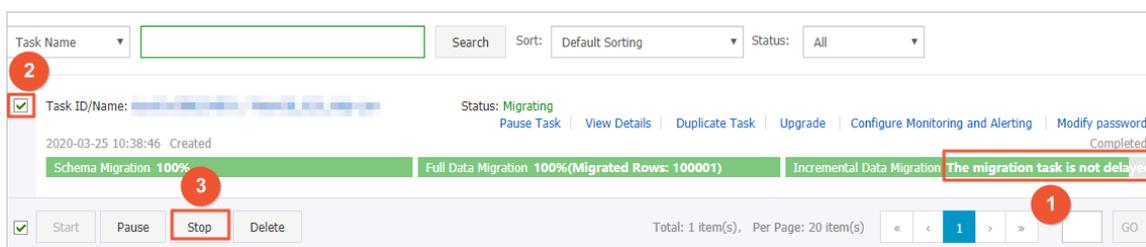
- Full data migration

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

- Incremental data migration

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

1. 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**.
2. After the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



6.13 Migrate incremental data from a user-created PostgreSQL database (version 9.4 to 9.6) 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 user-created PostgreSQL database to an ApsaraDB RDS for PostgreSQL instance](#).

Prerequisites

- The version of the user-created PostgreSQL database is 9.4, 9.5, or 9.6.
- 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 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.
- 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.
- During incremental data migration, only data manipulation language (DML) operations, such as INSERT, UPDATE, DELETE, and UPDATE operations, can be migrated.
- 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	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 #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

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

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 [#unique_73](#).

Data migration process	Description
1. Schema migration	DTS migrates the schemas of tables, views, sequences, functions, user-defined types, rules, domains, operations, and aggregates to the destination database. <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: Functions that are written in the C programming language cannot be migrated. </div>
2. Full data migration	DTS migrates historical data of the required objects to the destination database.
3. Schema migration	DTS migrates the schemas of triggers and foreign keys to the destination database.
4. Incremental data migration	DTS migrates incremental data of the required objects to the destination database. <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: Incremental data migration does not support the bit data type. </div>

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:
 - [Plug-in for PostgreSQL 9.4](#)
 - [Plug-in for PostgreSQL 9.5](#)
 - [Plug-in for PostgreSQL 9.6](#)

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';
```

- 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
```

- 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.

```
postgres=# SELECT * FROM pg_create_logical_replication_slot('replication_slot_test', 'ali_decoding');
 slot_name | xlog_position
-----+-----
 replication_slot_test | 0/5A
(1 row)
```

- 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 [#unique_62](#).

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

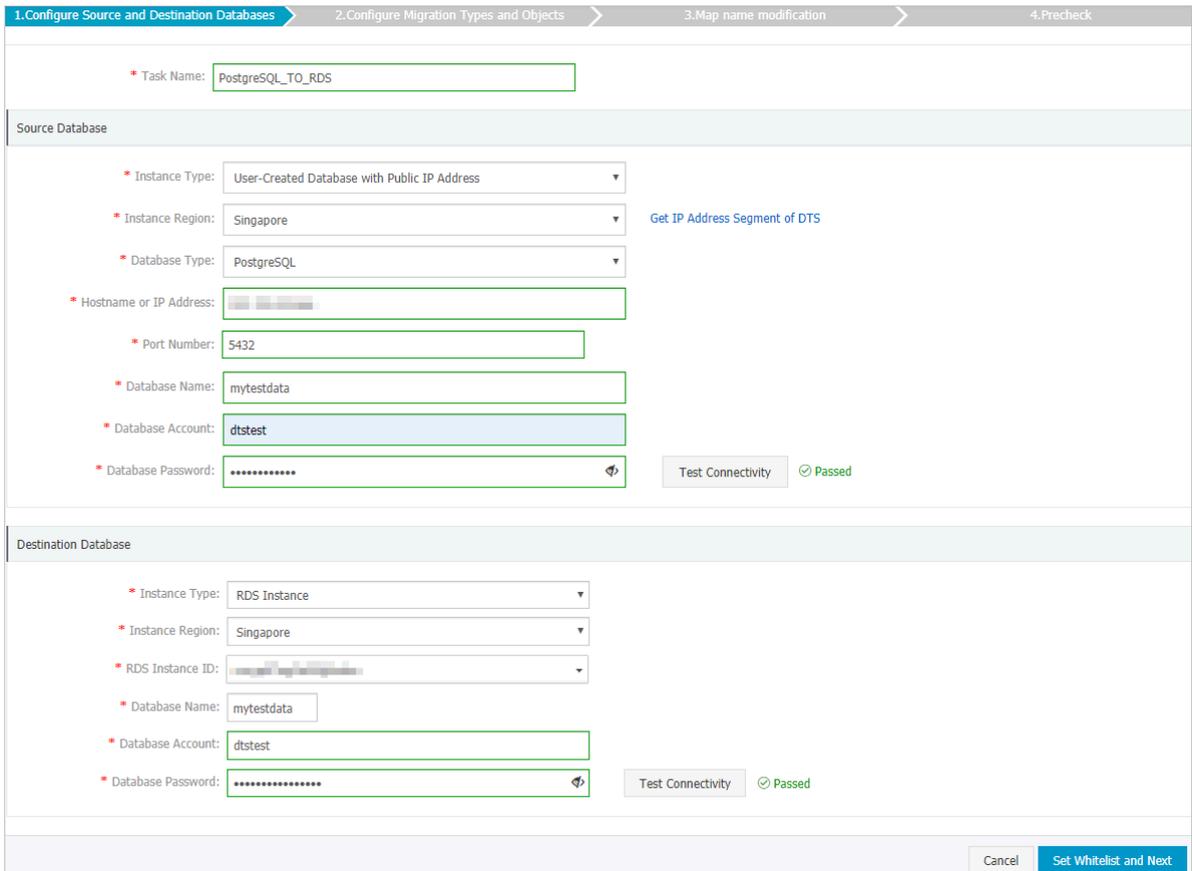
- Log on to the [DTS console](#).
- In the left-side navigation pane, click **Data Migration**.

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
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
Source Database	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 #unique_67 .
	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 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.
	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 .
	Database Name	Enter the name of the user-created PostgreSQL database.
	Database Account	Enter the account of the user-created PostgreSQL database. For more information about the permissions that are required for the account, see Permissions required for database accounts .

Section	Parameter	Description
	Database Password	<p>Enter the password for the source database account.</p> <p> 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.</p>
Destination Database	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 Name	<p>Enter the database name of the destination RDS instance. The name can be different from the name of the user-created PostgreSQL database.</p> <p> Note: Before you configure the data migration task, create a database and schema in the destination RDS instance. For more information, see Preparation 2: Create a database and schema in the destination RDS instance.</p>
	Database 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 .
	Database Password	<p>Enter the password for the destination database account.</p> <p> 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.</p>

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.

Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>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.</p> </div>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> You can select columns, tables, or databases as the objects to be migrated. You can select a database 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: <pre>ALTER TABLE schema.table REPLICA IDENTITY FULL;</pre> <p>Replace the schema and table in the preceding sample statement with the actual schema name and table name.</p> 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. 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**.



Note:

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

- Incremental data migration

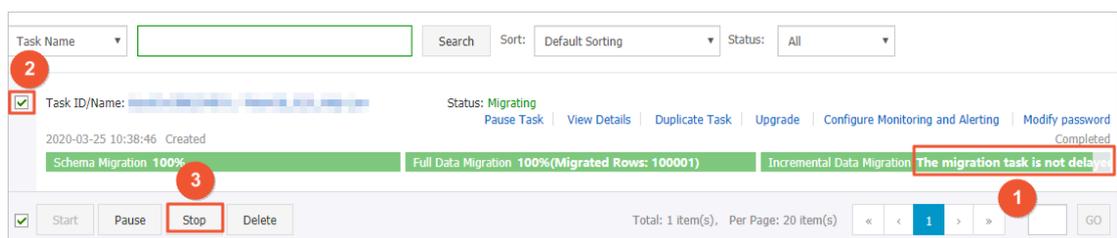
An incremental data migration task does not automatically stop. You must manually stop the migration task.



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.

- 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.
- 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 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.

6.14 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 user-created PostgreSQL database \(version 10.x to 12\) to an ApsaraDB RDS for PostgreSQL instance](#) and [Migrate incremental data from a user-created PostgreSQL database \(version 9.4 to 9.6\) 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 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.
- 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 #unique_51 .

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:

- [CREATE USER](#) and [GRANT](#) for a user-created PostgreSQL database
- [Create an account](#) for an ApsaraDB RDS for PostgreSQL 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. 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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type:

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

* Database Name:

* Database Account:

* Database Password: Passed

Destination Database

* Instance Type:

* Instance Region:

* RDS Instance ID:

* Database Name:

* Database Account:

* Database Password: Passed

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	<p>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.</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #f0f0f0; margin-top: 10px;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.</p> <p> 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.</p>
	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 .
	Database Name	Enter the name of the user-created PostgreSQL database.
	Database Account	Enter the account of the user-created PostgreSQL database. For more information about the permissions that are required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password for the source database account.</p> <p> 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.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination RDS instance resides.
	RDS Instance ID	Select the ID of the destination RDS instance.

Section	Parameter	Description
	Database Name	<p>Enter the database name of the destination RDS instance. The name can be different from the name of the user-created PostgreSQL database.</p> <p> Note: Before you configure the data migration task, create a database and schema in the destination RDS instance. For more information, see Preparations before data migration.</p>
	Database Account	<p>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.</p>
	Database Password	<p>Enter the password for the destination database account.</p> <p> 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.</p>

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.

The screenshot shows the '2.Configure Migration Types and Objects' step. At the top, there are four tabs: '1.Configure Source and Destination', '2.Configure Migration Types and Objects', '3.Map name modification', and '4.Precheck'. Under 'Migration Types', 'Schema Migration' and 'Full Data Migration' are checked, while 'Incremental Data Migration' is not. Below this, there is a warning about full data migration. The main area is split into two panels: 'Available' on the left and 'Selected' on the right. The 'Available' panel shows a tree view of database objects, with 'testschema' expanded to show 'Tables', 'Views', 'Sequences', 'Functions', 'User Defined Types', 'Rules', 'Domains', 'Operations', 'Aggregates', and 'Extensions'. The 'Selected' panel shows 'testschema (10 objects)' and 'customer'. At the bottom, there are buttons for 'Cancel', 'Previous', 'Save', and 'Precheck'.

Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <p>In this example, you must select Schema Migration and Full Data Migration.</p> <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p> Note: To ensure data consistency, we recommend that you do not write data into the user-created PostgreSQL database during full data migration.</p> </div>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> 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. </div>

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



Note:

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

6.15 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 [#unique_75](#).

- 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 noeviction, the data between the source and destination databases may become inconsistent. For more information about the data eviction policy, see [#unique_76](#)

- 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	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 #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

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
- BITOP, 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, MSETNX, and MULTI
- PERSIST, PEXPIRE, PEXPIREAT, PFADD, PFMERGE, PSETEX, and PUBLISH
- RENAME, RENAMENX, RESTORE, RPOP, RPOPLPUSH, RPUSH, and RPUSHX
- SADD, SDIFFSTORE, SELECT, SET, SETBIT, SETEX, SETNX, SETRANGE, SINTERSTORE, SMOVE, SPOP, SREM, and SUNIONSTORE
- 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 data 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.
- `<password>`: the password for the user-created Redis database.

Example

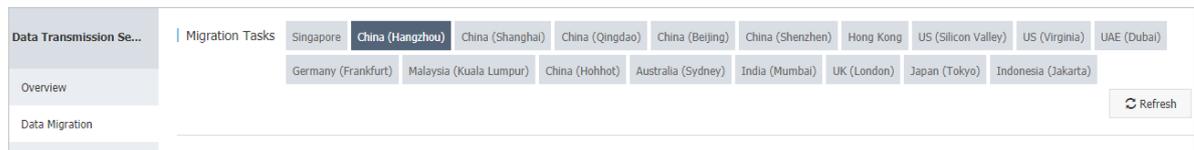
```
redis-cli -h 127.0.0.1 -p 6379 -a Test123456
```

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. 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.

1. Configure Source and Destination
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type:

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Instance Mode: Standalone

* Hostname or IP Address:

* Port Number:

Database Password:

✔ Passed

Destination Database

* Instance Type:

* Instance Region:

* Redis Instance ID:

Database Password:

✔ Passed

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	<p>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.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.</p> <p> 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.</p>
	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 .
	Database Password	<p>Enter the password for the user-created Redis database.</p> <p> 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.</p>
Destination Database	Instance Type	Select Redis Instance .
	Instance Region	Select the region where the destination ApsaraDB for Redis instance resides.
	Redis Instance ID	Select the ID of the destination ApsaraDB for Redis instance .

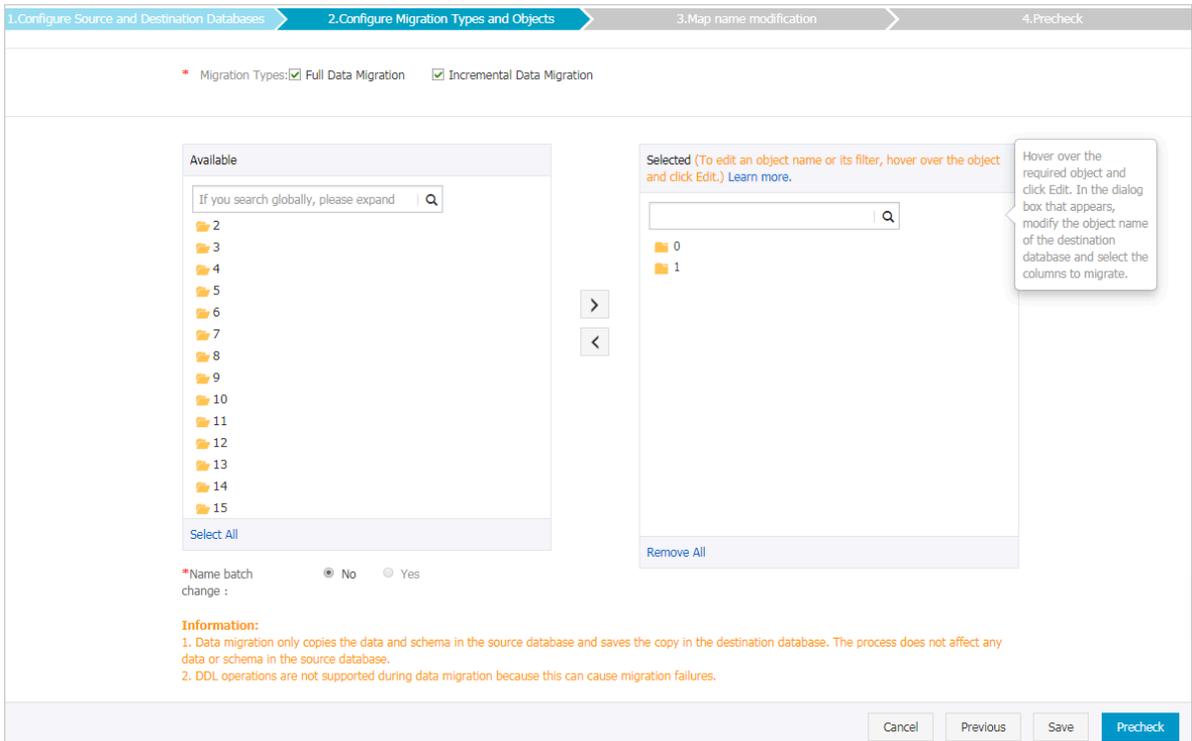
Section	Parameter	Description
	Database Password	<p>Enter the database password of the destination ApsaraDB for Redis instance.</p> <p> Note: After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the parameters are valid. If the specified 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.</p>

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 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.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <p>Note: If Incremental Data Migration is not selected, do not write data into the user-created Redis database during full data migration. This ensures data consistency between the source and destination databases.</p>
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p>Note: You can select databases as the objects to be migrated.</p>

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

Note:

- Before you can start the data migration task, a precheck is performed. You can start the data migration task only after the task passes the precheck.

- If the task fails to pass the precheck, 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.

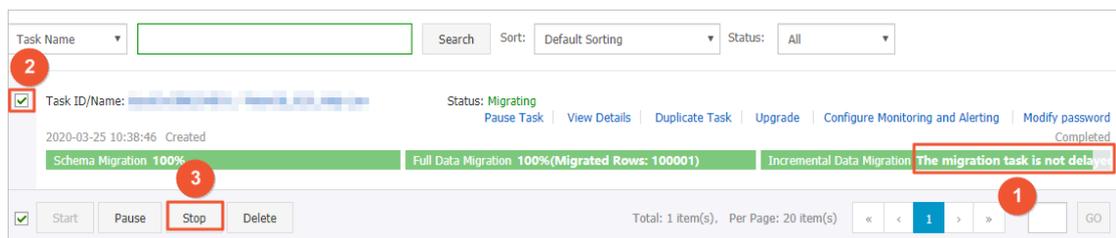
- Incremental data migration

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**, 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.

6.16 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 [#unique_77](#).

For more information about data migration or synchronization solutions, see [#unique_78](#).

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 [#unique_79](#).
- 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 Cloud over the Internet. For more information, see #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

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 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
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 through 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 MongoDB 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 --replSet 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.
- 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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Advanced Settings
4. Precheck

* Task Name:

Source Database

* Instance Type: [DTS support type](#)

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

Database Name: Authenticate Database with Account

Database Account:

Database Password: ✔ Passed

* Encryption: Non-encrypted SSL-encrypted(MongoDB Atlas only)

Destination Database

* Instance Type:

* Instance Region:

* MongoDB Instance ID:

* Database Name: Authenticate Database with Account

* Database Account:

* Database Password: ✔ Passed

Section	Parameter	Description
N/A	Task name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.

Section	Parameter	Description
Source Database	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 prepare the environments that are required for the source database. For more information, see #unique_67 .
	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.
	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.
	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.
	Database 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. For more information about the permissions that are required for the account, see Permissions required for database accounts .

Section	Parameter	Description
	Database Password	<p>Enter the password of the database account.</p> <p> 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 Failed message is displayed, you can click Check next to the Failed message to modify the information as prompted.</p>
	Connection Method	<p>Select Non-encrypted.</p> <p> Note: The SSL-encrypted option is available only when you migrate MongoDB Atlas.</p>
Destination Database	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.
	Database Name	<p>Enter the name of the authentication database to which the database account belongs.</p> <p> Note: If you want to use the root account, specify admin for the Database Name parameter.</p>
	Database Account	<p>Enter the username of the database account that you use to manage the destination database. For more information about the permissions that are required for the account, see Permissions required for database accounts.</p>
	Database Password	<p>Enter the password of the database account.</p> <p> 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.</p>

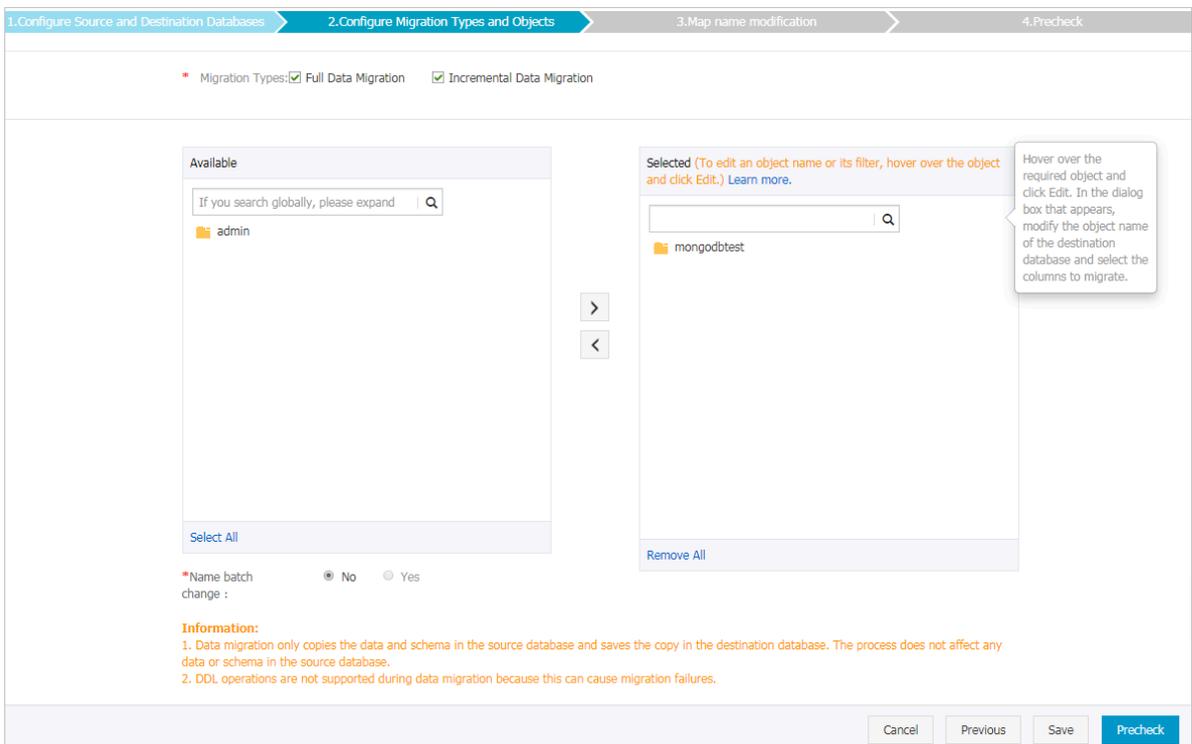
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p>Note:</p> <ul style="list-style-type: none"> 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. </div>

Parameter	Description
Migration objects	<ul style="list-style-type: none"> Select objects from the Available section and click the  icon to move the objects to the Selected section. <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p> Note:</p> <ul style="list-style-type: none"> - 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. </div> <ul style="list-style-type: none"> 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**.



Note:

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

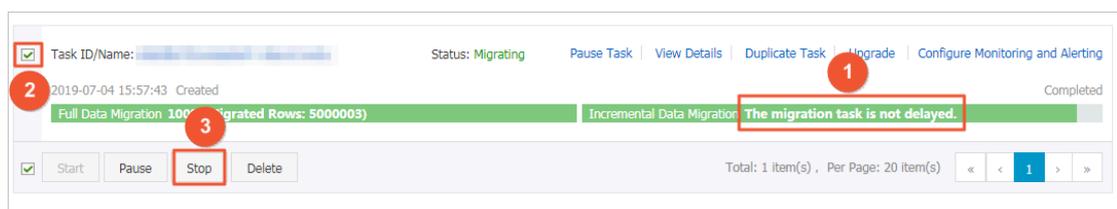
- Incremental data migration

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.

- 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.
- 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.**

6.17 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 [#unique_82](#).

For more information about data migration or synchronization solutions, see [#unique_78](#).

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 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 [#unique_79](#).

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, visit #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Advanced Settings
4. Precheck

* Task Name:

Source Database

* Instance Type: DTS support type

* Instance Region: Get IP Address Segment of DTS

* Database Type:

* Hostname or IP Address:

* Port Number:

Database Name: Authenticate Database with Account

Database Account:

Database Password: Test Connectivity ✔ Passed

* Encryption: Non-encrypted SSL-encrypted(MongoDB Atlas only)

Destination Database

* Instance Type:

* Instance Region:

* MongoDB Instance ID:

* Database Name: Authenticate Database with Account

* Database Account:

* Database Password: Test Connectivity ✔ Passed

Cancel Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	<p>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.</p> <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If Instance Type is set to User-Created Database with Public IP Address, you do not need to specify the Instance Region parameter.</p> <p> 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.</p>
	Database 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	<p>Enter the service port of the user-created MongoDB database.</p> <p> Note: In this example, the service port of the user-created MongoDB database must be open to the public network.</p>
	Database Name	Enter the name of the authentication database to which the database account belongs.
	Database Account	Enter the account that is used to log on to the user-created MongoDB database. For more information about the permissions that are required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password of the source database account.</p> <p> 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.</p>

Section	Parameter	Description
	Connection Method	Select Non-encrypted .  Note: The SSL-encrypted option is available only when you migrate MongoDB Atlas.
Destination Database	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 database.
	Database Name	Enter the name of the authentication database to which the database account belongs.  Note: If you want to use the root account, specify admin for the Database Name parameter.
	Database Account	Enter the account that is used to log on to the ApsaraDB for MongoDB 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 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 **Set Whitelist and Next**.

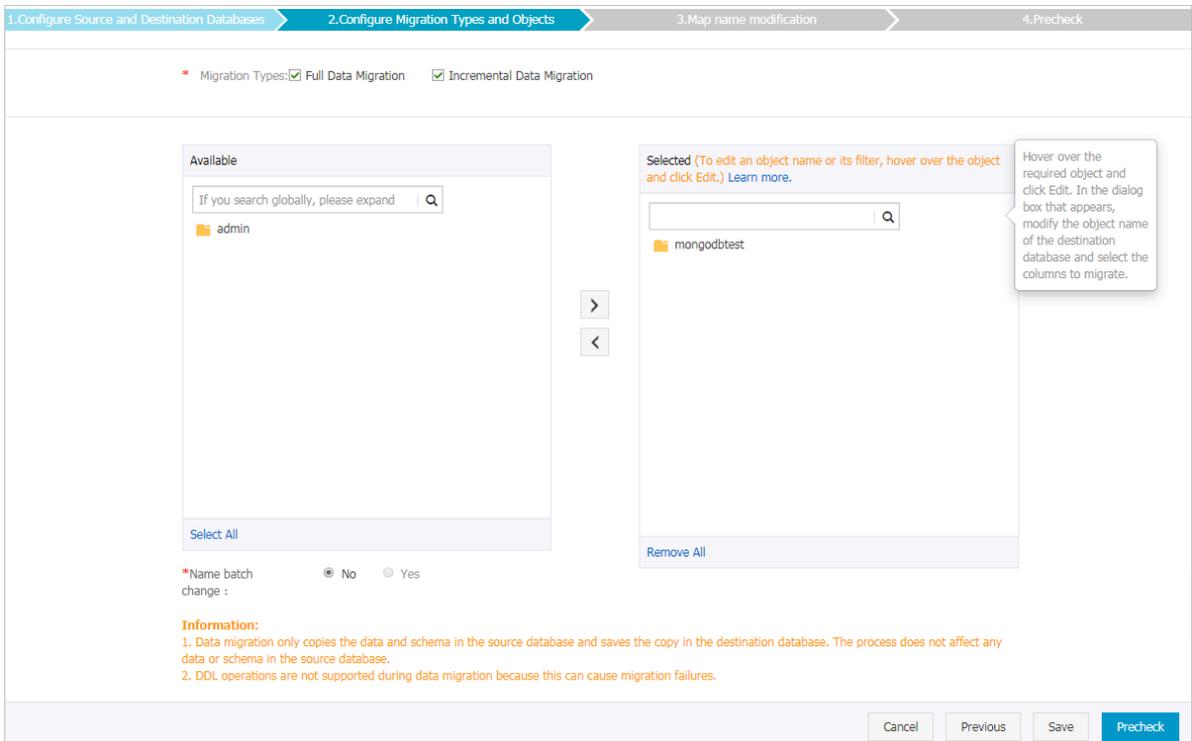


Note:

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 an ApsaraDB for MongoDB instance.](#)

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



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <p> Note: 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.</p>

Parameter	Description
Migration objects	<ul style="list-style-type: none"> Select objects from the Available section and click the  icon to move the objects to the Selected section. <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p> Note:</p> <ul style="list-style-type: none"> - 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. </div> <ul style="list-style-type: none"> 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**.



Note:

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

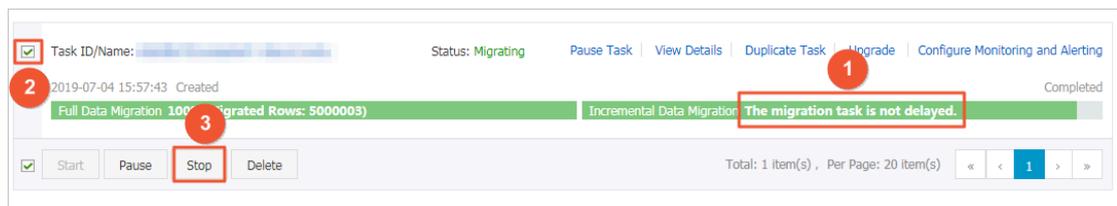
- Incremental data migration

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.

- 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.
- 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?](#)

6.18 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.

To avoid service disruptions, we recommend that you use DTS to migrate user-created MongoDB databases to ApsaraDB for MongoDB. You can also use the built-in commands of MongoDB to migrate user-created MongoDB databases. For more information, see [#unique_84](#).

For more information about data migration and synchronization solutions, see [#unique_78](#).

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.

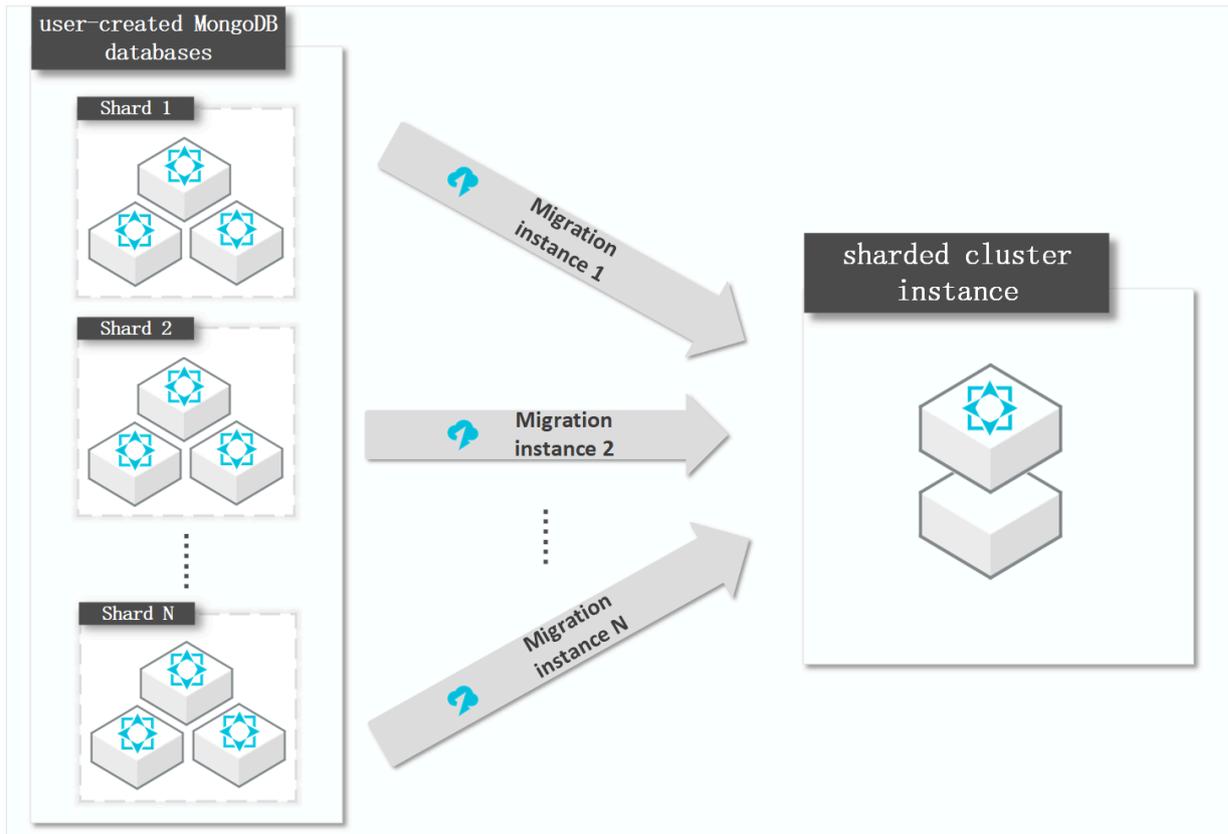
How it works

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 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 [#unique_79](#).

Billing

Migration type	Instance configurations	Internet traffic
Full data migration	Free of charge.	Charged only when data is migrated from Alibaba Cloud over the public network. For more information, visit #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

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 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.

- a. Download the [cleanupOrphaned.js](#) file.

```
wget "http://docs.aliyun.cn-hangzhou.oss.aliyun-inc.com/assets/attach/120562/cn_zh/1564451237979/cleanupOrphaned.js"
```

- b. 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.

```
function cleanupOrphaned(coll) {
  var nextKey = { };
  var result;

  while ( nextKey != null ) {
    result = db.adminCommand( { cleanupOrphaned: coll, startingFromKey: nextKey } );

    if (result.ok != 1)
      print("Unable to complete at this time: failure or timeout.")

    printjson(result);

    nextKey = result.stoppedAtKey;
  }
}

var dbName = 'test'
db = db.getSiblingDB(dbName)
db.getCollectionNames().forEach(function(collName) {
  cleanupOrphaned(dbName + "." + collName);
});
```

- c. 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 <passwd> 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 'Test123456' cleanupOrphaned.js
```

```
mongo --host 172.16.1.11 --port 27021 --authenticationDatabase admin -u root -p 'Test123456' cleanupOrphaned.js
```

```
mongo --host 172.16.1.12 --port 27024 --authenticationDatabase admin -u root -p 'Test123456' 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 **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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Advanced Settings
4. Precheck

* Task Name:

Source Database

* Instance Type: DTS support type

* Instance Region: Get IP Address Segment of DTS

* Database Type:

* Hostname or IP Address:

* Port Number:

Database Name: Authenticate Database with Account

Database Account:

Database Password: Test Connectivity ✔ Passed

* Encryption: Non-encrypted SSL-encrypted(MongoDB Atlas only)

Destination Database

* Instance Type:

* Instance Region:

* MongoDB Instance ID:

* Database Name: Authenticate Database with Account

* Database Account:

* Database Password: Test Connectivity ✔ Passed

Cancel Set Whitelist and Next

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	<p>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.</p> <div style="background-color: #f2f2f2; padding: 10px; margin-top: 10px;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If Instance Type is set to User-Created Database with Public IP Address, you do not need to specify the Instance Region.</p> <p> 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.</p>
	Database Type	Select MongoDB .
	Hostname or IP Address	<p>Enter the endpoint of a shard for the source database. In this example, enter the public IP address of the shard.</p> <p> Note: 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 shard in the second migration task. Repeat this operation until all shards are migrated.</p>
	Port Number	<p>Enter the service port of the shard.</p> <p> Note: The service port of each shard for user-created MongoDB databases must be open to the public network.</p>
	Database 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. For more information about the permissions that are required for the account, see Permissions required for database accounts .

Section	Parameter	Description
	Database Password	<p>Enter the password of the source database account.</p> <p> 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 Failed message is displayed, click Check in the Failed message to modify the information as prompted.</p>
	Connection Method	<p>Select Non-encrypted.</p> <p> Note: The SSL-encrypted option is available only when you migrate MongoDB Atlas.</p>
Destination Database	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	<p>Enter the name of the authentication database to which the database account belongs.</p> <p> Note: If you want to use the root account, specify admin for the Database Name parameter.</p>
	Database Account	<p>Enter the username of the database account that you use to manage the source database. For more information about the permissions that are required for the account, see Permissions required for database accounts.</p>

Section	Parameter	Description
	Database Password	<p>Enter the password for the destination database account.</p> <p> 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.</p>

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.

1. Configure Source and Destination 2. Configure Migration Types and Objects 3. Map name modification 4. Precheck

Migration Types: Full Data Migration Incremental Data Migration

Available

If you search globally, please expand

- admin

Select All

Selected (To edit an object name or its filter, hover over the object and click Edit.) [Learn more.](#)

- mongodbtest

Remove All

*Name batch change : No Yes

Information:

1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.
2. DDL operations are not supported during data migration because this can cause migration failures.

Cancel Previous Save **Precheck**

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



Note:

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



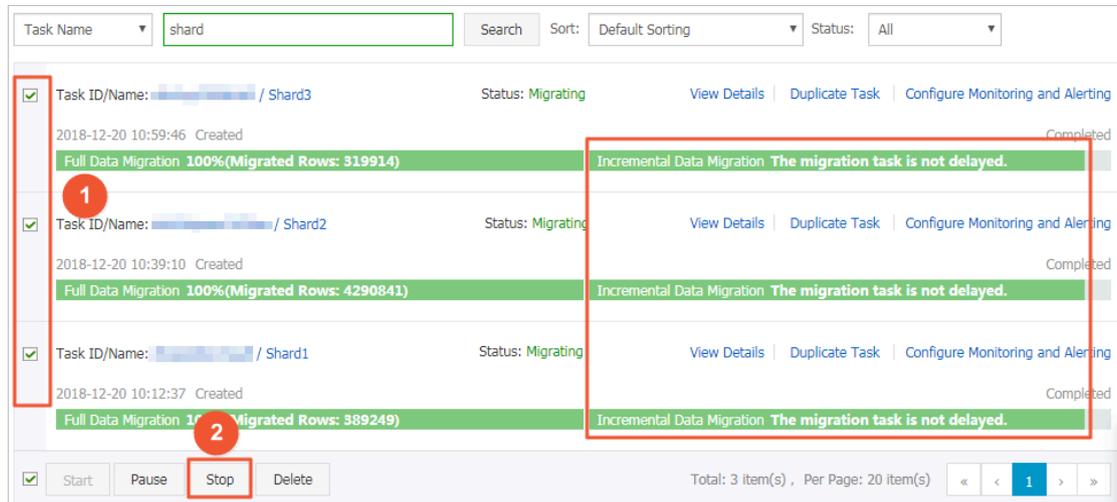
Note:

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.

6.19 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

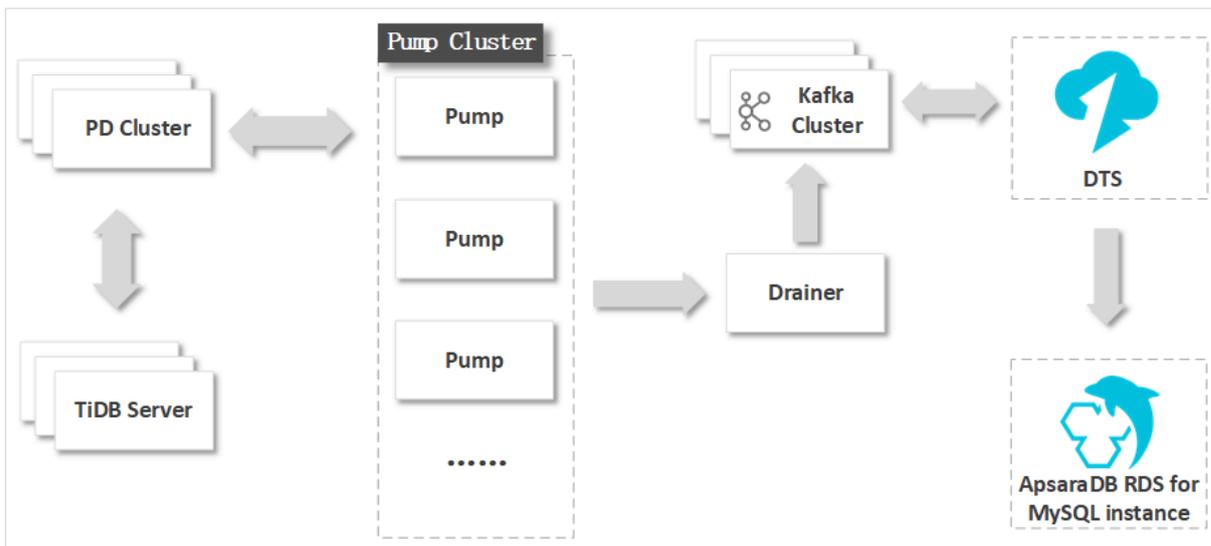
An ApsaraDB RDS for MySQL instance is created. For more information, see [#unique_65](#).



Notice:

- The destination ApsaraDB RDS for MySQL instance resides in the China (Hangzhou), China (Shanghai), China (Qingdao), China (Beijing), China (Shenzhen), China (Zhangjiakou-Beijing Winter Olympics), China (Hong Kong), Singapore, US (Silicon Valley), or 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.

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 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

Migration type	Description
Schema migration	<p>DTS migrates the schemas of the required objects to the destination database. DTS supports schema migration for views, tables, and databases.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> 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.</p> </div>

Migration type	Description
Full data migration	<p>DTS migrates historical data of the required objects to the destination database.</p> <p> 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.</p>
Incremental data migration	<p>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:</p> <ul style="list-style-type: none"> • Data manipulation language (DML) operations: INSERT, UPDATE, and DELETE • Data definition language (DDL) operations: CREATE TABLE, DROP TABLE, ALTER TABLE, RENAME TABLE, TRUNCATE TABLE, CREATE VIEW, DROP VIEW, and ALTER VIEW <p>Incremental data migration allows you to ensure service continuity when you migrate data from a user-created TiDB database to Alibaba Cloud.</p>

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 [#unique_62](#).

Procedure

1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the 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.

a) Configure the task name and source database.

* Task Name:

Note: If you need to carry out incremental data migration for a long time, it is recommended to use the data synchronization function, which has better network stability and richer functions. [Click to buy data synchronization](#) [Click to see the difference between data migration and synchronization](#)

Source Database

* Instance Type: DTS support type

* Instance Region:

* ECS Instance ID:

* Database Type:

* Port Number:

* Database Account:

* Database Password:

* Incremental migration or not: Yes No

* Kafka ClusterType:

* Instance Region:

* ECS Instance ID:

* KafkaPort Number:

Kafka Cluster Account: Optional

Kafka Cluster Password: Optional

* Topic:
Click Get Topic List and then select the specific topic.

* Kafka Version:

* Kafka ClusterEncryption: Non-encrypted SCRAM-SHA-256

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.

Parameter	Description
Instance Type	<p>Select an instance type based on where the source database is deployed.</p> <p>User-Created Database in ECS Instance is used in this example.</p> <div data-bbox="470 790 651 1485"> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</div>
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.

Parameter	Description
Database Password	<p data-bbox="459 271 632 472">Enter the password of the source database account.</p> <div data-bbox="459 495 652 2027"> Notice: 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.</div>

Parameter	Description
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 user-created TiDB database to an ApsaraDB RDS for MySQL instance.

Parameter	Description
Kafka Cluster Type	<p>Select an instance type based on where the user-created Kafka cluster is deployed.</p> <p>User-Created Database with Public IP Address is selected in this example.</p> <div data-bbox="459 824 644 1541" style="background-color: #f0f0f0; padding: 5px;">  Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67. </div>
Instance Region	The value of this parameter is the same as the region of the source database and cannot be changed.

Parameter	Description
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.

Parameter	Description
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.

b) Configure the destination database.

Destination Database

- * Instance Type:
- * Instance Region:
- * RDS Instance ID:
- * Database Account:
- * Database Password:

* Encryption: Non-encrypted SSL-encrypted

Parameter	Description
Instance Type	Select RDS Instance .
Instance Region	Select the region where the destination RDS instance resides.

Parameter	Description
Database Account	Enter the database account of the destination RDS instance. The account must have the read/write permissions on the destination database. For more information about how to create and authorize a database account, see Create an account for an ApsaraDB RDS for MySQL instance and Modify the permissions of an account for an RDS instance .

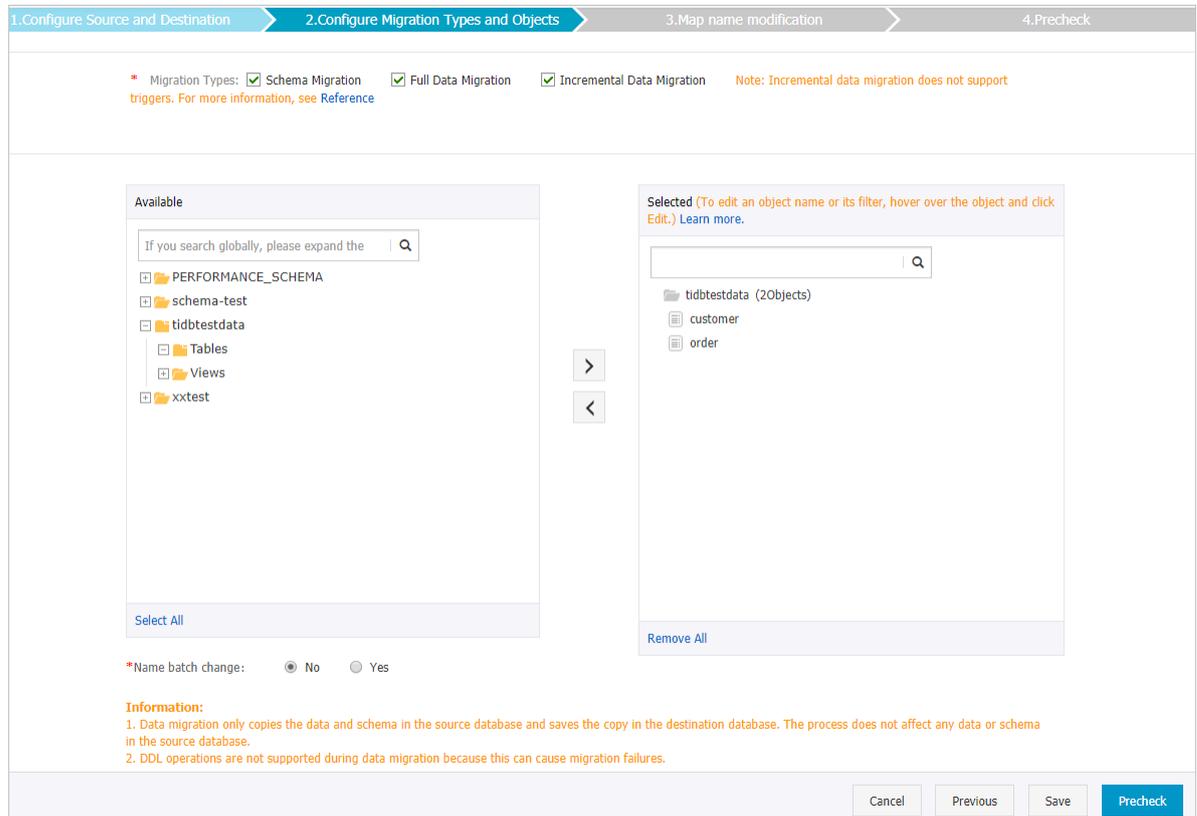
Parameter	Description
Database Password	<p>Enter the password of the destination database account.</p> <div data-bbox="467 533 652 2074" style="background-color: #f0f0f0; padding: 5px;">  Notice: After you specify the source database parameters, click Test Connectivity next to Database Password to verify 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. </div>

Parameter	Description
Encryption	<p>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.</p> <p> Notice: The Encryption parameter is available only for regions in mainland China and the China (Hong Kong) region.</p>

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.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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, all of the three migration types are selected.

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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.

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

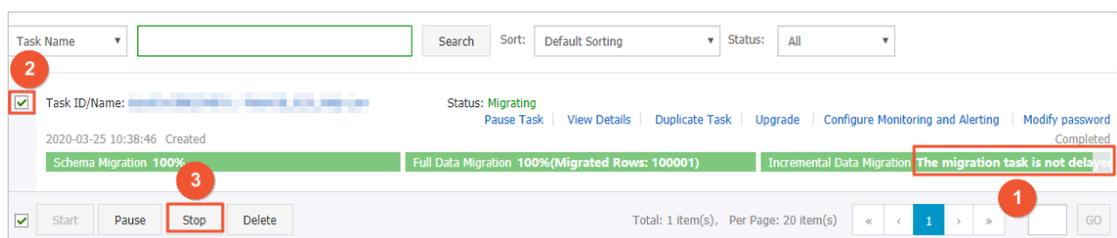
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



6.20 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 [#unique_65](#).



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 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	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 #unique_51 .

Migration types

DTS reads binary logs of the source database to migrate incremental data. The implementation mechanism of binary logs in TiDB is different from that in ApsaraDB RDS for MySQL. Therefore, DTS does not support incremental data migration from a TiDB database to an ApsaraDB RDS for MySQL instance.

- Schema migration

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



Notice:

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](#).

- Full data migration

DTS migrates historical data of the required objects to the destination database.



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.
- To ensure data consistency, do not write data to the source database during full data migration.

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 [#unique_62](#).

Procedure

1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the 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.

The screenshot shows the configuration interface for a data migration task. It is divided into four steps: 1. Configure Source and Destination, 2. Configure Migration Types and Objects, 3. Advanced Settings, and 4. Precheck. The current step is 1. The 'Task Name' is 'TiDB_To_RDS'. The 'Source Database' section includes fields for Instance Type (User-Created Database in ECS Instance), Instance Region (China (Hangzhou)), ECS Instance ID (i-bp1-...), Database Type (TiDB), Port Number (4000), Database Account (dtstest), and Database Password (masked). A 'Test Connectivity' button shows 'Passed'. The 'Destination Database' section includes fields for Instance Type (RDS Instance), Instance Region (China (Hangzhou)), RDS Instance ID (rm-bp-...), Database Account (dtstest), and Database Password (masked). A 'Test Connectivity' button shows 'Passed'. There are radio buttons for 'Encryption' with 'Non-encrypted' selected. At the bottom right, there are 'Cancel' and 'Set Whitelist and Next' buttons.

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.
Source Database	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. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;">  Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67. </div>
	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 .

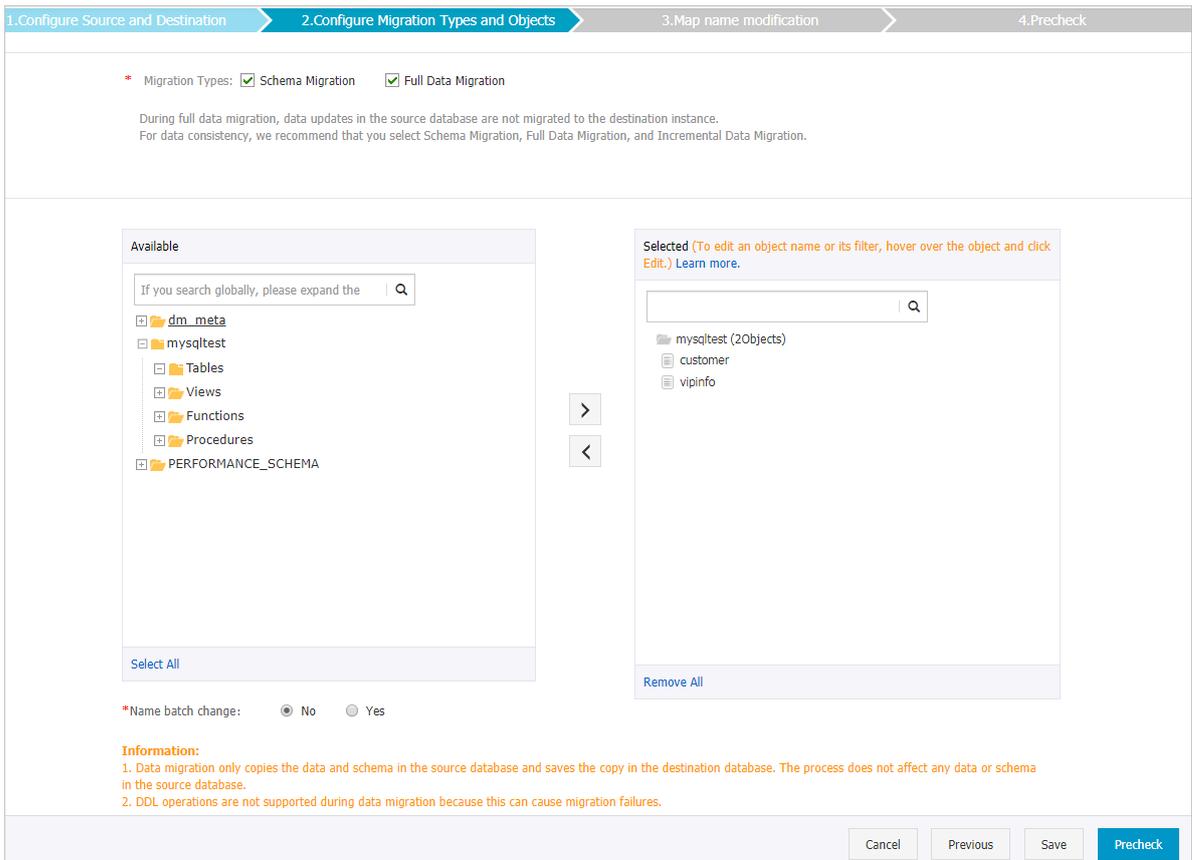
Section	Parameter	Description
	Database Account	Enter the account of the source TiDB database. For more information about the permissions that are required for the account, see Permissions required for database accounts .
	Database Password	Enter the password of the source database account.  Notice: 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.
Destination Database	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 the permissions that are required for the account, see Permissions required for database accounts .
	Database Password	Enter the password of the destination database account.  Notice: 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.

Section	Parameter	Description
	Encryption	<p>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.</p> <p> Notice: The Encryption parameter is available only for regions in mainland China and the Hong Kong (China) region.</p>

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.



Parameter	Description
Migration Types	<p>Select both Schema Migration and Full Data Migration. For more information about migration types, see Migration types.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Notice: To ensure data consistency, do not write data to the source TiDB database during data migration.</p> </div>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Notice:</p> <ul style="list-style-type: none"> 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. </div>

8. Click **Precheck** on the lower right of the page.

 **Note:**

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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.

Task ID/Name: XXXXXXXXXX Status: Completed [View Details](#) | [Duplicate Task](#) | [Configure Monitoring and Alerting](#)

2019-05-24 17:01:02 Created 2019-05-27 15:54:33 Completed

Schema Migration 100%

Full Data Migration 100%(Migrated Rows: 1000000)

Start Pause Stop Delete Total: 1 Item(s), Per Page: 20 Item(s) 1

6.21 Migrate data from a user-created Db2 database to an ApsaraDB RDS for MySQL database

This topic describes how to migrate data from a user-created Db2 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. To 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 database is larger than the total space of the data in the user-created Db2 database.

Notes

- DDL operations cannot be synchronized.
- If the name of the source database is invalid, you must create a database in the ApsaraDB RDS 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 databases and accounts](#).

- 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 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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

Migration types

- Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, index, and foreign key.

- Full data migration

DTS migrates historical data of the required objects from the source Db2 database to the destination ApsaraDB RDS for MySQL database.

- Incremental data migration

After full data migration is complete, DTS synchronizes incremental data from the user-created Db2 database to the destination ApsaraDB RDS for MySQL database. Incremental data migration helps you 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
User-created Db2 database	The CONNECT and SELECT permissions	The CONNECT and SELECT permissions	The DBADM permission
ApsaraDB RDS for MySQL database	The read/write permissions	The read/write permissions	The read/write permissions

For more 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 RDS for MySQL instance](#) and [Change the permissions of an account for an RDS for MySQL instance](#) for an ApsaraDB RDS for MySQL database

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 incremental data migration.

Preparations before incremental data migration

Before configuring an incremental data migration task, you must enable the archive log feature of the source Db2 database. For more information, see [Primary log archive method](#) and [Secondary log archive method](#).



Note:

Skip this step if you only perform full 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 destination RDS instance 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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type:

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

* Database Name:

* Database Account:

* Database Password:

Test Connectivity ✔ Passed

Destination Database

* Instance Type:

* Instance Region:

* RDS Instance ID:

* Database Account:

* Database Password:

Test Connectivity ✔ Passed

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 Database	Instance Type	<p>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.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>

Section	Parameter	Description
	Instance Region	<p>If the instance type is set to User-Created Database with Public IP Address, you do not need to specify the instance region.</p> <p> 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 user-created Db2 database. You can click Get IP Address Segment of DTS next to Instance Region to obtain the CIDR blocks of DTS servers.</p>
	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.
	Port Number	Enter the service port number of the user-created Db2 database. The default port number is 50000 .
	Database Name	Enter the name of the source Db2 database.
	Database Account	Enter the account for the user-created Db2 database. For more information about permissions required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password for the database account.</p> <p> Note: 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 as prompted.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination RDS instance resides.
	RDS Instance ID	Select the ID of the destination RDS instance.

Section	Parameter	Description
	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 .
	Database Password	Enter the password for the database account.  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 as prompted.
	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).

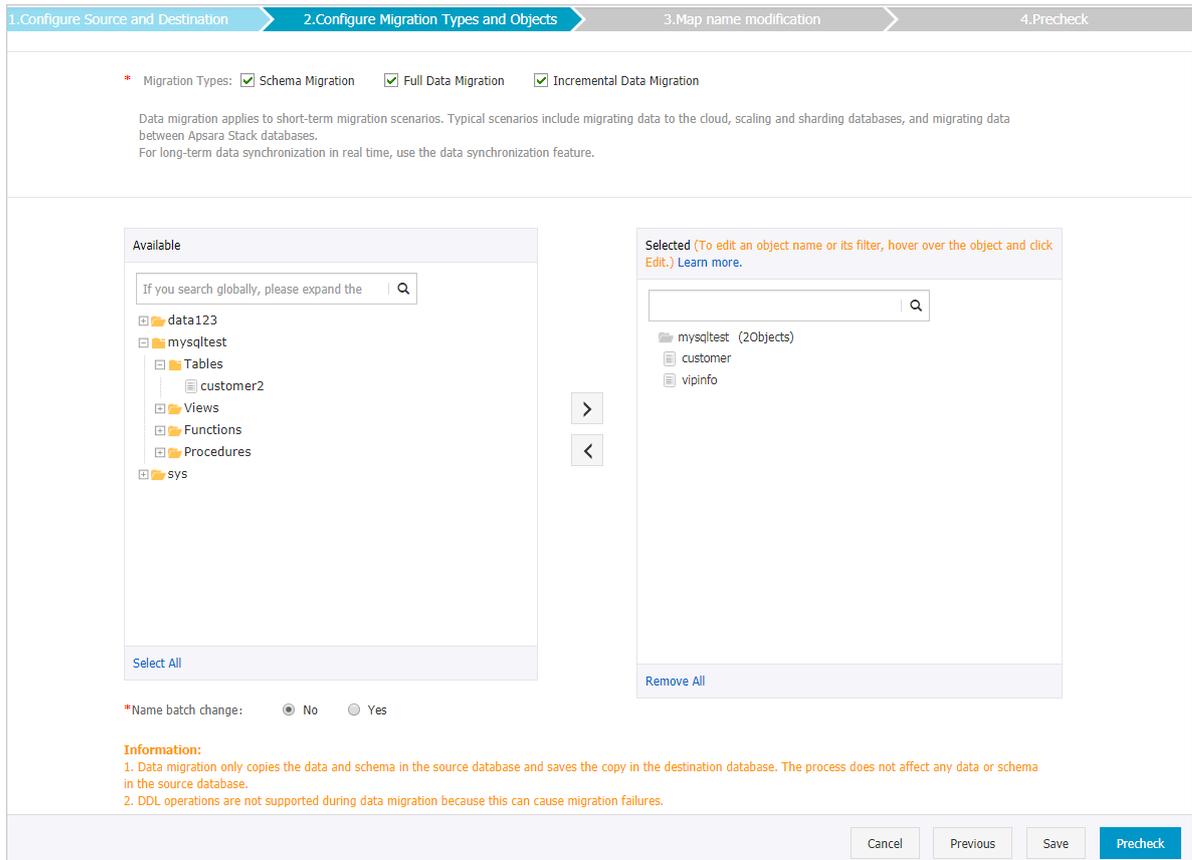
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



Note:

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

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



Parameter	Description
Migration types	<ul style="list-style-type: none"> 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. <p>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.</p>

Parameter	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> You can select databases, tables, or columns as the objects to be migrated. After a Db2 database is migrated to the destination RDS instance, the name of the Db2 database remains unchanged. If you want an object to have a different name after the object is migrated to the destination RDS 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 on an object, other objects that are dependent on the object may fail to be migrated.

8. Click **Precheck** in the lower-right corner of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  corresponding to each failed item to view the details.

Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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 migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Incremental data migration

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



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination instance.

- When the task progress bar switches to **Incremental Data Migration** and **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.

The screenshot shows the Data Migration console interface. At the top, there is a search bar and filters for Task Name, Sort (Default Sorting), and Status (All). Below this, a task is listed with a checkbox selected. The task ID/name is partially visible. The status is 'Migrating'. There are links for 'Pause Task', 'View Details', 'Duplicate Task', 'Upgrade', 'Configure Monitoring and Alerting', and 'Modify password'. The progress bar shows three stages: 'Schema Migration 100%' (green), 'Full Data Migration 100%(Migrated Rows: 100001)' (green), and 'Incremental Data Migration' (green) with a sub-status 'The migration task is not delayed' (red box). At the bottom, there are buttons for 'Start', 'Pause', 'Stop' (highlighted with a red box and a red circle '3'), and 'Delete'. The bottom right shows pagination: 'Total: 1 item(s), Per Page: 20 item(s)' and a page number '1' (highlighted with a red circle '1') and a 'GO' button.

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

7 Migration from a third-party cloud to Alibaba Cloud

7.1 Migrate data from an Amazon RDS for MySQL database to an ApsaraDB RDS for MySQL database

This topic describes how to migrate data from an Amazon RDS for 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 RDS for MySQL is set to **Yes**. This ensures that DTS can access Amazon RDS for MySQL over the Internet.
- An ApsaraDB RDS for MySQL instance is created. For more information, see [Create an RDS for MySQL instance](#).
- The available storage space of the destination ApsaraDB RDS for MySQL database is larger than the total space of the data in the Amazon RDS for 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

- Schema migration

DTS migrates the schemas of the required objects to the destination instance. DTS supports schema migration for the following types of objects: table, view, trigger, stored procedure, and function. DTS does not support schema migration for events.



Note:

- During schema migration, DTS changes the value of the SECURITY attribute from DEFINER to INVOKER for views, stored procedures, and functions.

- DTS does not migrate user information. 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 source Amazon RDS for 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.
- To ensure data consistency, do not write new data into the Amazon RDS for MySQL instance during full data migration.

- Incremental data migration

After full data migration is complete, DTS retrieves binary log files from the source Amazon RDS for MySQL database. Then, DTS synchronizes incremental data from the source Amazon RDS for 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

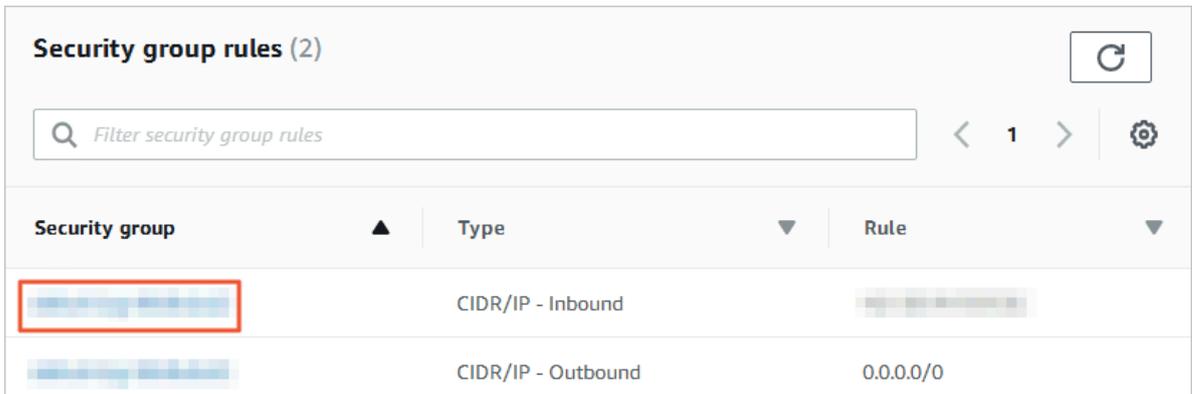
Database	Schema migration	Full data migration	Incremental data migration
Amazon RDS for MySQL database	The SELECT permission	The SELECT permission	The SELECT, REPLICATION SLAVE, and REPLICATION CLIENT permissions
ApsaraDB RDS for MySQL database	The read/write permissions	The read/write permissions	The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

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

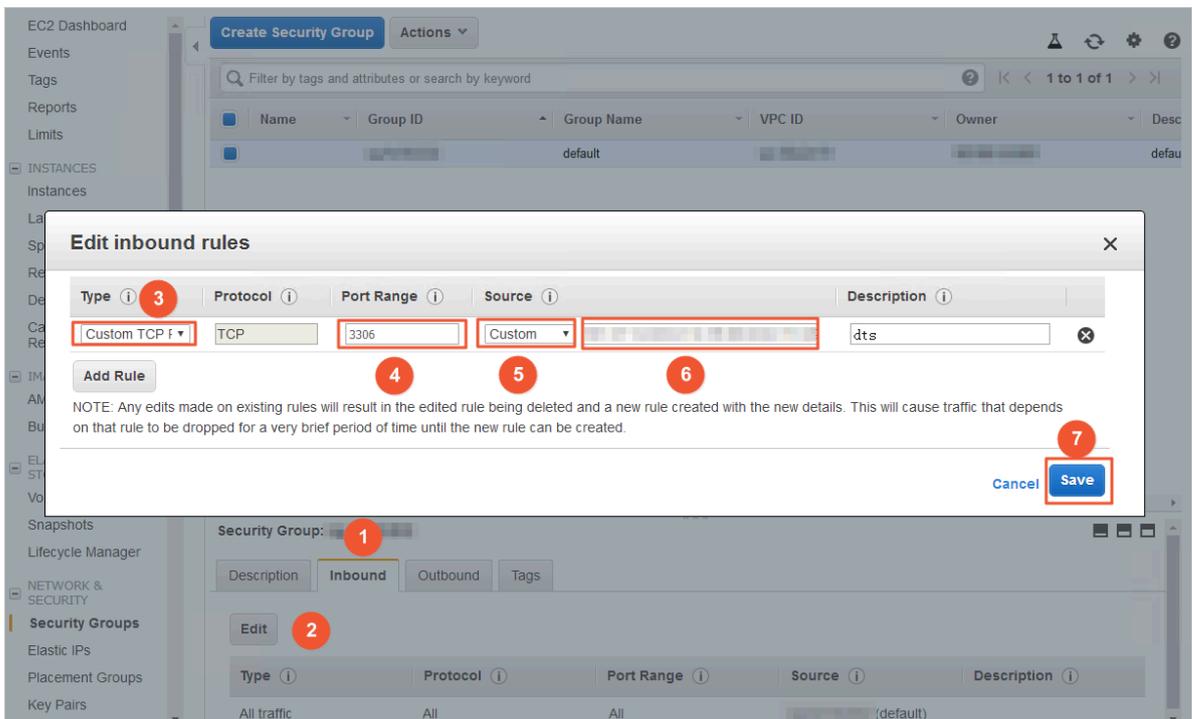
- [#unique_66](#) for an Amazon RDS for MySQL database
- [Create an account for an RDS for MySQL instance](#) and [Change the permissions of an account for an RDS for MySQL instance](#) for an ApsaraDB RDS for MySQL database

Preparations before data migration

1. Log on to the Amazon RDS Management Console.
2. Go to the **Basic information** page of the source 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 CIDR blocks of DTS servers in the corresponding region to the inbound rule. For more information, see [#unique_62](#).



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.

5. Log on to the Amazon RDS for MySQL database and specify the number of hours to retain binary log files. Skip this step if you do not need to perform incremental data migration.

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

**Note:**

- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of Amazon RDS for 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. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.

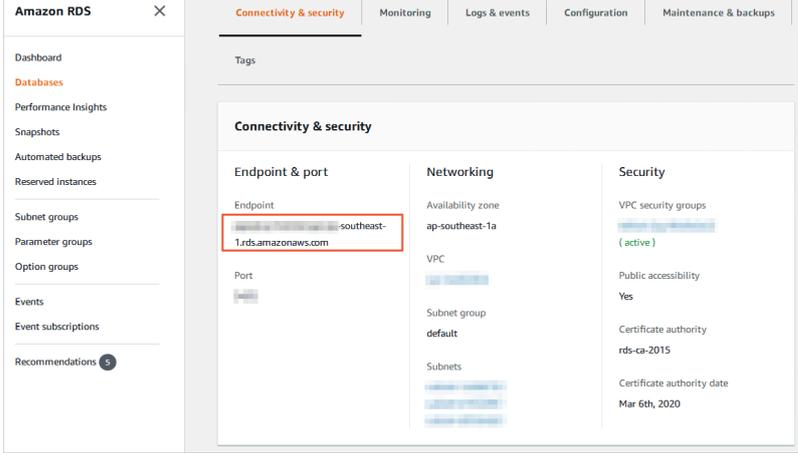


4. In the upper-right corner of the page, click **Create Migration Task**.

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

The screenshot shows the '1. Configure Source and Destination' step of a DTS migration task. The task name is 'MySQL_TO_MySQL'. The source database is configured as a 'User-Created Database with Public IP Address' in the 'Singapore' region, with a 'MySQL' database type. The port number is 3306, and the database account is 'dtstest'. The destination database is an 'RDS Instance' in the 'Singapore' region, with a database account of 'dtstest'. Both source and destination connectivity tests have passed.

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 Database	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
	Hostname or IP Address	<p>Enter the endpoint that is used to connect to the Amazon RDS for MySQL database.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  <p>Note: You can obtain the endpoint on the Basic information page of the source Amazon RDS for MySQL instance.</p> </div> 
	Port Number	Enter the service port number of the Amazon RDS for MySQL database. The default port number is 3306 .
	Database Account	Enter the account for the Amazon RDS for MySQL database. For more information about permissions required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password for the database account.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  <p>Note: 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 as prompted.</p> </div>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the ApsaraDB RDS for MySQL instance resides.
	RDS Instance ID	Select the ID of the ApsaraDB RDS for MySQL instance.

Section	Parameter	Description
	Database 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 .
	Database Password	Enter the password for the database account.  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 as prompted.
	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).

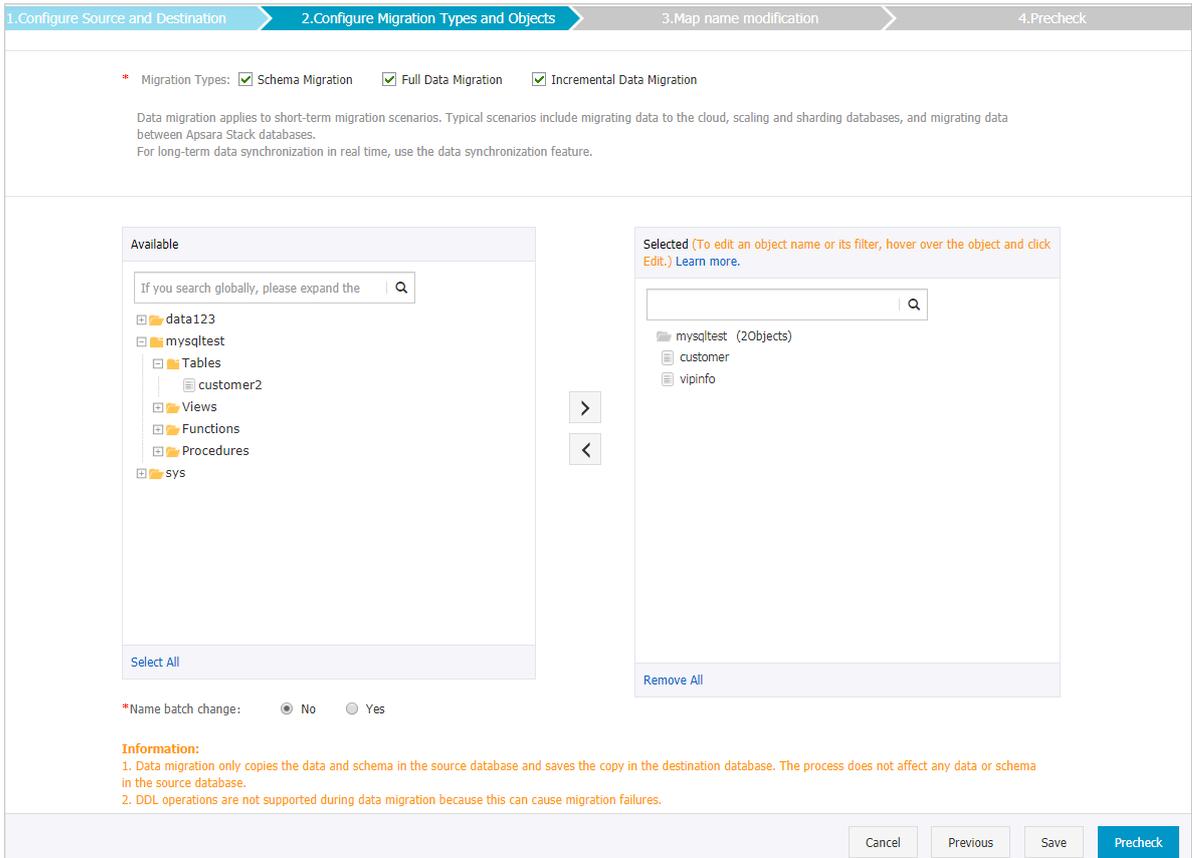
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



Note:

The CIDR blocks of DTS servers are automatically added to the whitelist of the 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.



Parameter	Description
Migration types	<ul style="list-style-type: none"> 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. <p>Note:</p> <ul style="list-style-type: none"> 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. 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.

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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.

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

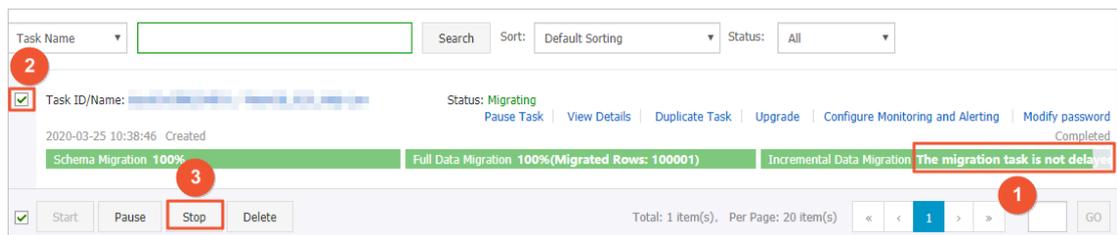
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- When 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 MySQL instance.

7.2 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 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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 data 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.

Security group	Type	Rule
[Red Box]	CIDR/IP - Inbound	[Redacted]
[Redacted]	CIDR/IP - Outbound	0.0.0.0/0

4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit, 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 [#unique_62](#).

Edit inbound rules

Type	Protocol	Port Range	Source	Description
Custom TCP f	TCP	1521	Custom	dts

NOTE: Any edits made on existing rules will result in the edited rule being deleted and a new rule created with the new details. This will cause traffic that depends on that rule to be dropped for a very brief period of time until the new rule can be created.

Buttons: Add Rule, Cancel, Save



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.

5. Change log configurations for the Amazon RDS for Oracle database. Skip this step if you do not need to perform incremental data migration.

- a) Use the master user account and the SQL*Plus tool provided by Oracle to connect to the Amazon RDS for Oracle database.
- b) 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](#).

- c) Enable force logging.

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

- d) Enable supplemental logging for primary keys.

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

- e) Enable supplemental logging for unique keys.

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

- f) Set a retention period for archived logs.

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

**Note:**

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

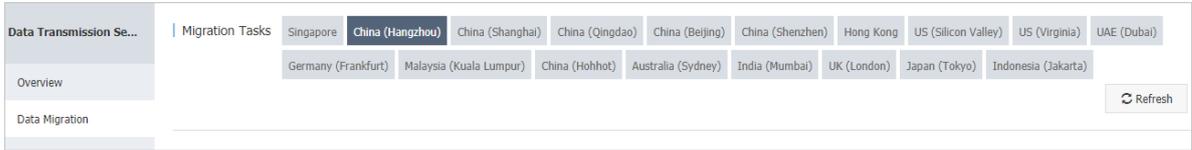
- g) Commit the changes.

```
commit;
```

Procedure

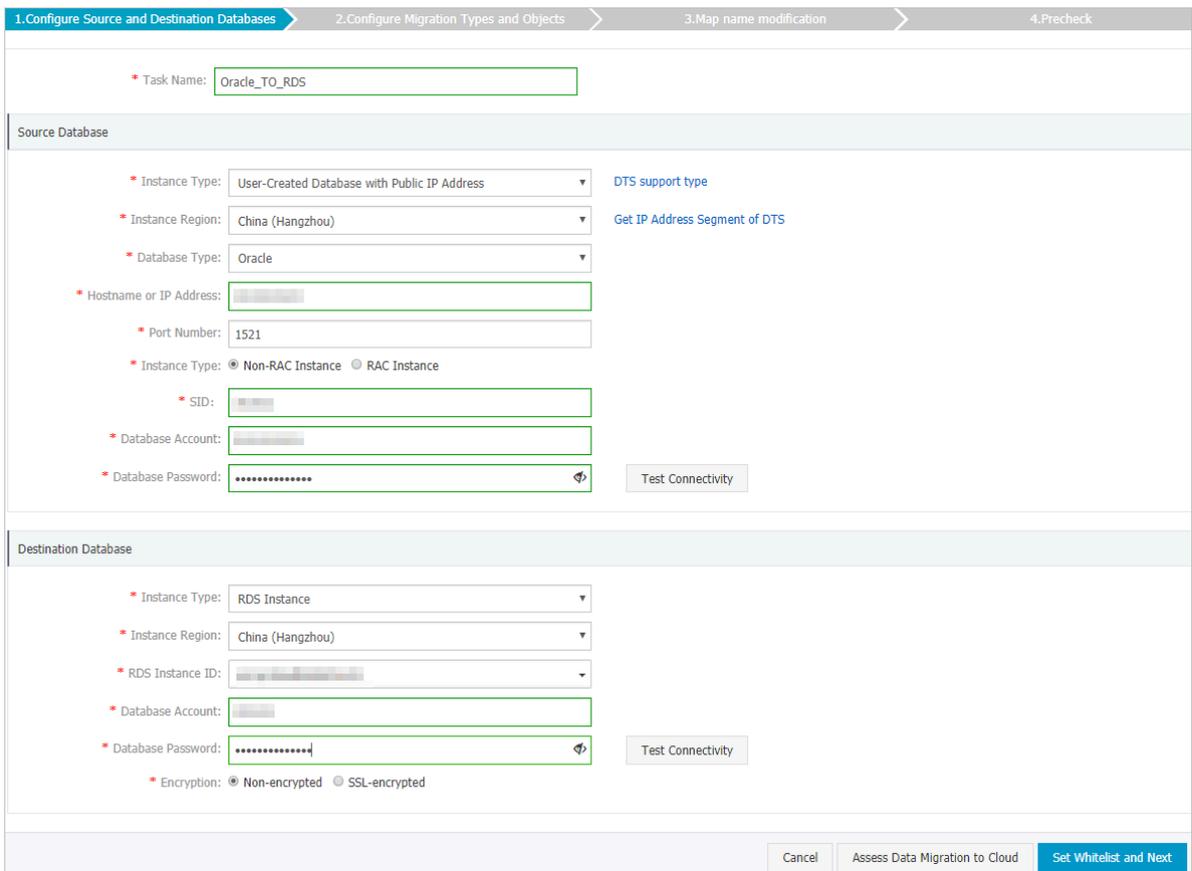
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.

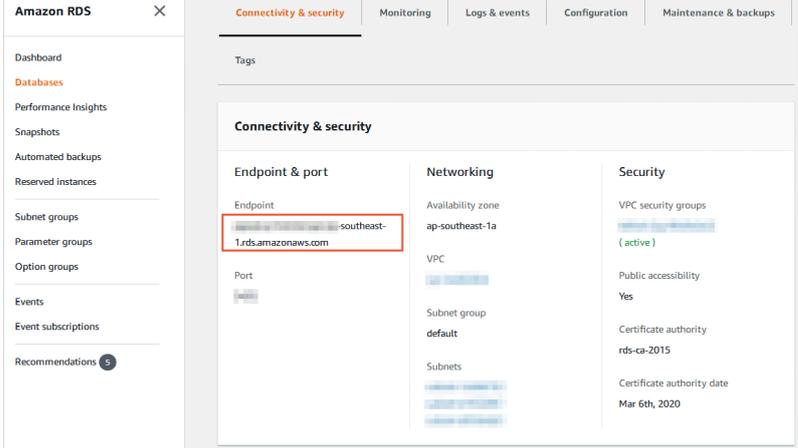


4. In the upper-right corner of the page, click **Create Migration Task**.

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.
Source Database	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
	Hostname or IP Address	<p>Enter the endpoint for accessing the Amazon RDS for Oracle database.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  <p>Note: You can obtain the endpoint on the Basic information page of the source Amazon RDS for Oracle database.</p> </div> 
	Port Number	<p>Enter the port number configured for the Amazon RDS for Oracle database. The default port number is 1521.</p>
	Instance Type	<ul style="list-style-type: none"> • 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. <p>In this example, select Non-RAC Instance and enter SID.</p>
	Database Account	<p>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.</p>
	Database Password	<p>Enter the password of the database account.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  <p>Note: 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 as prompted.</p> </div>

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 RDS instance ID.
	Database 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 .
	Database Password	Enter the password of the database account.  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 as prompted.

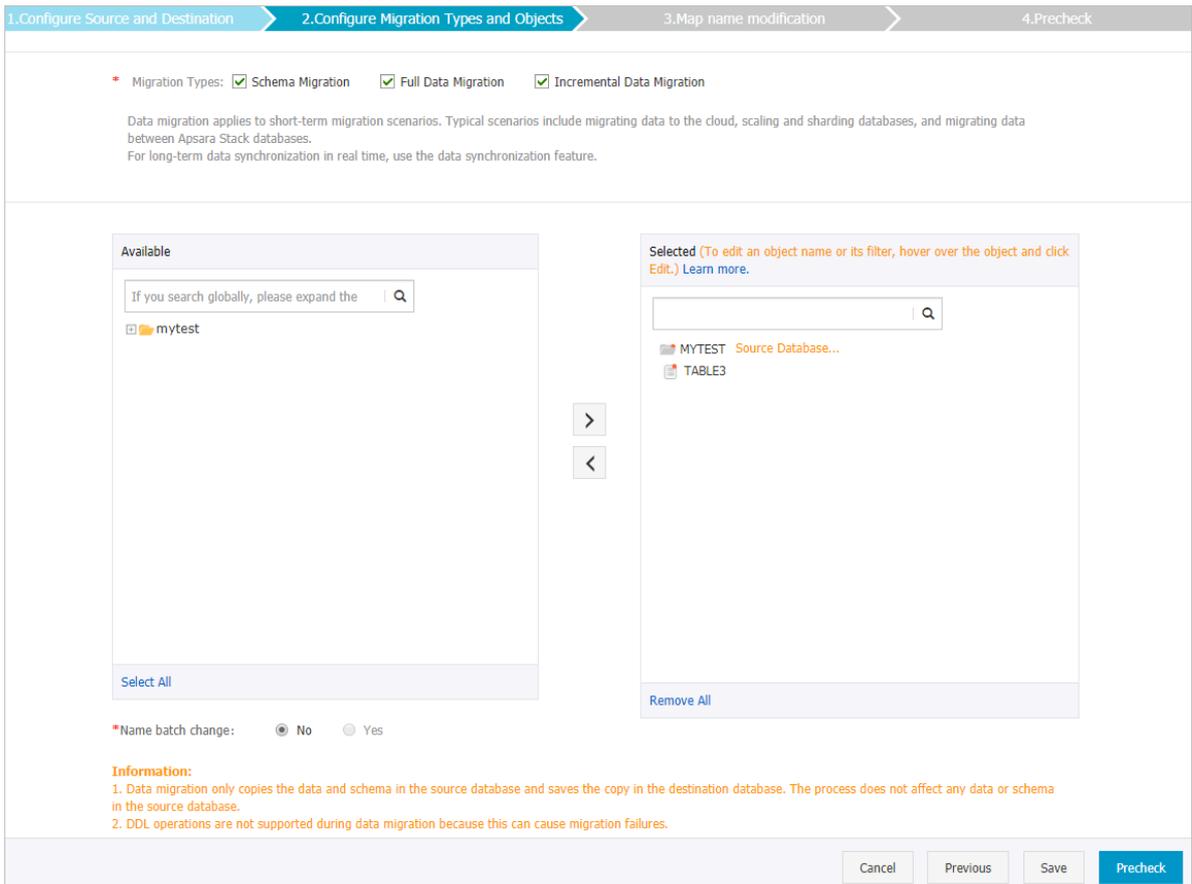
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



Note:

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

7. Configure migration types and objects.



Section	Description
Migration Types	<ul style="list-style-type: none"> To perform only full data migration, select Schema Migration and Full Data Migration. To achieve service continuity during migration, select Schema Migration, Full Data Migration, and Incremental Data Migration. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Section	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

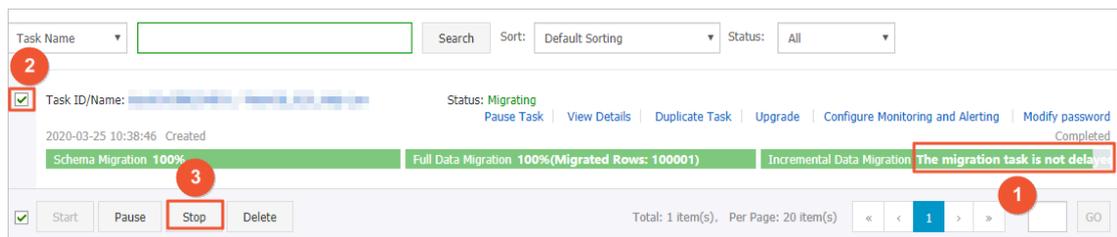
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When 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 MySQL database.

7.3 Migrate full data from an Amazon RDS for Oracle database to an ApsaraDB RDS for PPAS database

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

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.
- An ApsaraDB RDS for PPAS instance is created. For more information, see [Create an ApsaraDB RDS for PPAS instance](#).
- The available storage space of the ApsaraDB RDS for PPAS database is larger than the total size of the data in the Amazon RDS for Oracle database.

Background information

- 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.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the task . Otherwise, the data in the source database will overwrite the data in the destination database 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 public network incurs fees. For more information, see #unique_51 .

Permissions required for database accounts

Database	Schema migration	Full data migration	Incremental data migration
Amazon RDS for Oracle	Owner permissions on schemas to be migrated	Owner permissions on schemas to be migrated	Master user permissions
ApsaraDB RDS for PPAS database	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 and authorize database account, see the following topics:

- [CREATE USER](#) and [GRANT](#) for an Amazon RDS for Oracle database
- [Create an account for an RDS for PPAS instance](#)

Limits

- Incremental data migration is not supported, because the dbcreator and sysadmin roles required for incremental data migration are not available in Amazon RDS for Oracle.
- Do not perform DDL operations on the source database during full data migration because these operations cannot be synchronized.
- Migration of materialized views is not supported.
- During schema migration, reverse indexes and bitmap indexes are converted into common indexes. Partitioned indexes become indexes in each partitioned table of the ApsaraDB RDS for PPAS database.

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.

Security group	Type	Rule
[Red Boxed]	CIDR/IP - Inbound	[Redacted]
[Redacted]	CIDR/IP - Outbound	0.0.0.0/0

4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit, 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 [#unique_62](#).

Edit inbound rules

Type: Custom TCP f (3) | Protocol: TCP (4) | Port Range: 1521 (5) | Source: Custom (6) | Description: dts

NOTE: Any edits made on existing rules will result in the edited rule being deleted and a new rule created with the new details. This will cause traffic that depends on that rule to be dropped for a very brief period of time until the new rule can be created.

Buttons: Add Rule (4), Cancel, Save (7)

Background: Security Group: [Redacted] (1) | Inbound tab | Edit button (2)



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

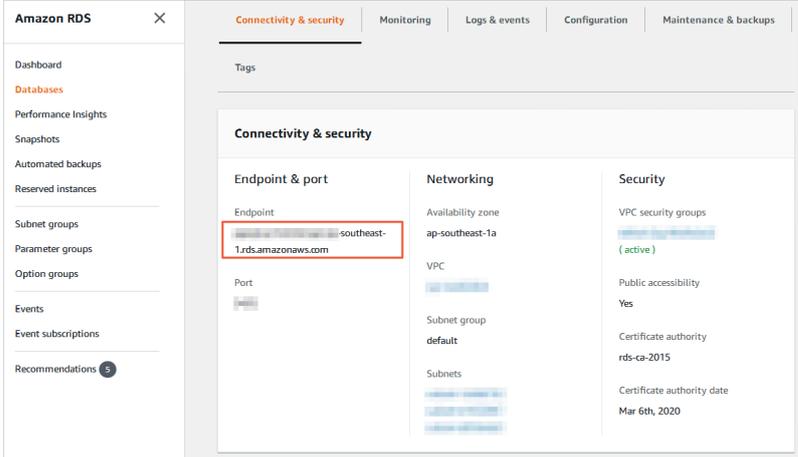
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.



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

The screenshot shows the '1. Configure Source and Destination Databases' step of the migration task creation process. The interface is divided into four steps: 1. Configure Source and Destination Databases (active), 2. Configure Migration Types and Objects, 3. Map name modification, and 4. Precheck. The 'Task Name' field is set to 'Oracle_TO_PPAS'. The 'Source Database' section includes fields for Instance Type (User-Created Database with Public IP Address), Instance Region (Singapore), Database Type (Oracle), Hostname or IP Address, Port Number (1521), Instance Type (Non-RAC Instance selected), SID, Database Account (dtstest), and Database Password. A 'Test Connectivity' button is present. The 'Destination Database' section includes fields for Instance Type (RDS Instance), Instance Region (Singapore), RDS Instance ID, Database Name (mytest), Database Account (dtstest), and Database Password. Another 'Test Connectivity' button is present. At the bottom right, there are 'Cancel' and 'Set Whitelist and Next' buttons.

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.

Section	Parameter	Description
Source Database	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 .
	Hostname or IP Address	Enter the endpoint for accessing the Amazon RDS for Oracle database. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;">  Note: You can obtain the endpoint on the Basic information page of the source Amazon RDS for MySQL instance. </div> 
	Port Number	Enter the port number configured for the user-created Oracle database. The default port number is 1521 .
	Instance Type	<ul style="list-style-type: none"> • 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.
	Database Account	Enter the account for the Amazon RDS for MySQL database. For more information about permissions required for the account, see Permissions required for database accounts .

Section	Parameter	Description
	Database Password	<p>Enter the password of the database account.</p> <p> Note: 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 as prompted.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination ApsaraDB RDS for PPAS instance resides.
	RDS Instance ID	Select the destination ApsaraDB RDS for PPAS instance ID.
	Database Name	Enter the name of the destination database.
	Database Account	Enter the database account of the destination ApsaraDB RDS for PPAS instance. For more information about permissions required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password of the database account.</p> <p> 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 as prompted.</p>

6. Click **Set Whitelist and Next** in the lower-right corner of the page.



Note:

The IP addresses of DTS servers are automatically added to the whitelist of the destination RDS instance. Then, DTS servers can connect to the RDS instance.

7. Configure migration types and objects.

1.Configure Source and Destination
2.Configure Migration Types and Objects
3.Map name modification
4.Precheck

* Migration Types: Schema Migration Full Data Migration Incremental Data Migration

Data migration applies to short-term migration scenarios. Typical scenarios include migrating data to the cloud, scaling and sharding databases, and migrating data between Apsara Stack databases.
For long-term data synchronization in real time, use the data synchronization feature.

Available

If you search globally, please expand the

mytest

[Select All](#)

Selected (To edit an object name or its filter, hover over the object and click Edit.) [Learn more.](#)

MYTEST Source Database...

TABLE3

[Remove All](#)

>
<

*Name batch change: No Yes

Information:

1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.
2. DDL operations are not supported during data migration because this can cause migration failures.

Cancel Previous Save Precheck

Item	Description
Migration types	<p>Select Schema Migration and Full Data Migration.</p> <div style="background-color: #f2f2f2; padding: 10px; margin-top: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> Incremental data migration is not supported because the dbcreator and sysadmin roles required for incremental data migration are not available in Amazon RDS for Oracle. To ensure data consistency, do not write new data into the Amazon RDS for MySQL instance during full data migration. </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.

 **Note:**

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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 migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

Task ID/Name: [redacted] Status: Completed [View Details](#) | [Duplicate Task](#) | [Configure Monitoring and Alerting](#)

2019-05-24 17:01:02 Created 2019-05-27 15:54:33 Completed

Schema Migration 100%

Full Data Migration 100%(Migrated Rows: 1000000)

Start Pause Stop Delete Total: 1 item(s), Per Page: 20 item(s) « < 1 > »

12. Switch your workloads to the destination ApsaraDB RRDS for PPAS database.

7.4 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](#).



Note:

- 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 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 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.



Note:

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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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.

- Incremental 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.

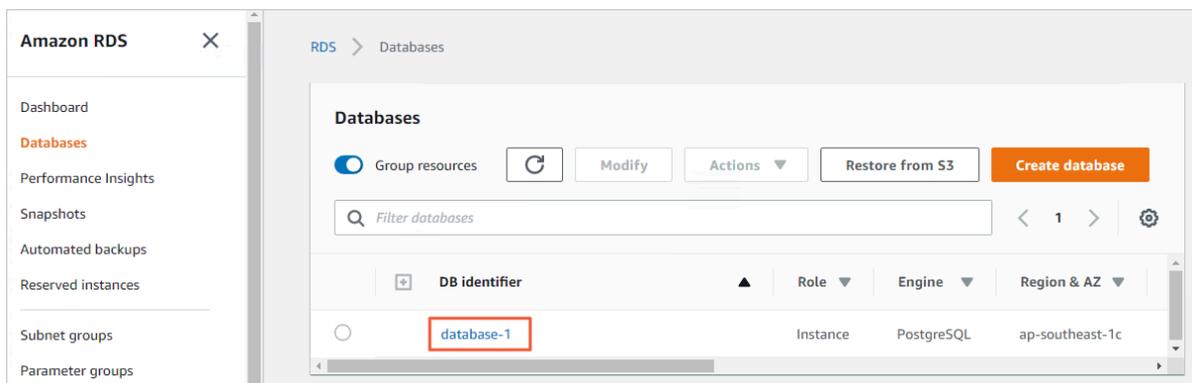
2. Perform full data migration.
3. Migrate the schemas of triggers and foreign keys.
4. Perform incremental data migration.

**Note:**

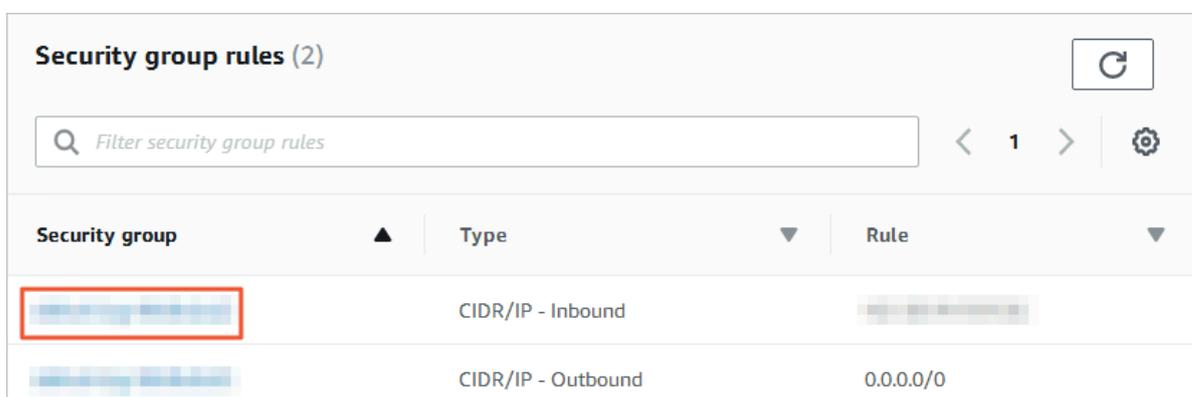
Before performing incremental data migration, do not perform DDL operations on the objects to be migrated in the Amazon RDS for PostgreSQL database. Otherwise, the objects may fail to be migrated.

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.

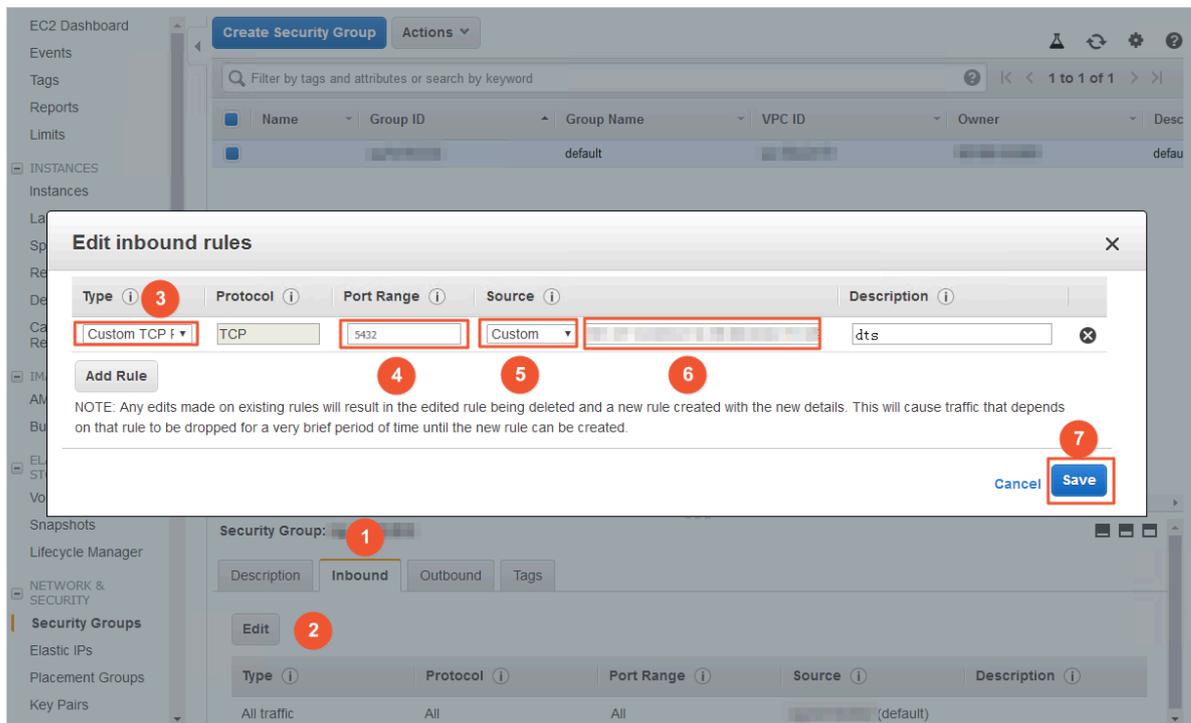


4. In the **Security group rules** section, click the name of the security group corresponding to the existing inbound rule.



5. 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 [#unique_62](#).



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

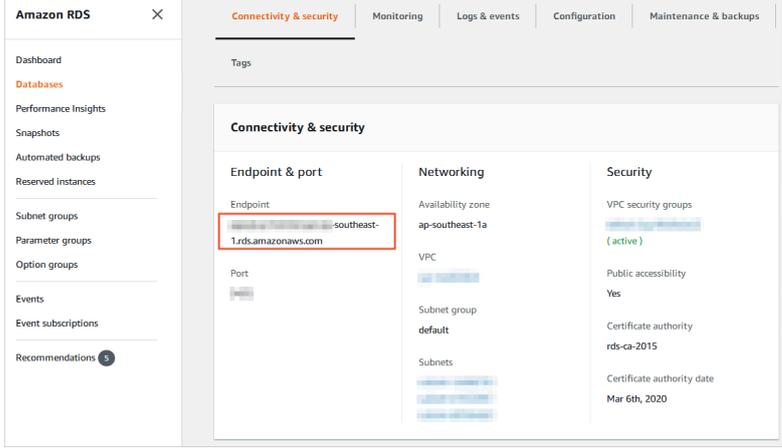
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.



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.
Source Database	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
	Hostname or IP Address	<p>Enter the endpoint that is used to connect to the Amazon RDS for PostgreSQL database.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  Note: You can find the endpoint on the Basic information page of the source Amazon RDS for PostgreSQL instance. </div> 
	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.
	Database 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 .
	Database Password	<p>Enter the password for the database account.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  Note: 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 as prompted. </div>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.

Section	Parameter	Description
	RDS Instance ID	Select the ID of the destination ApsaraDB RDS for PostgreSQL instance.
	Database 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.  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 .
	Database 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 .
	Database Password	Enter the password for the database account.  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 as prompted.

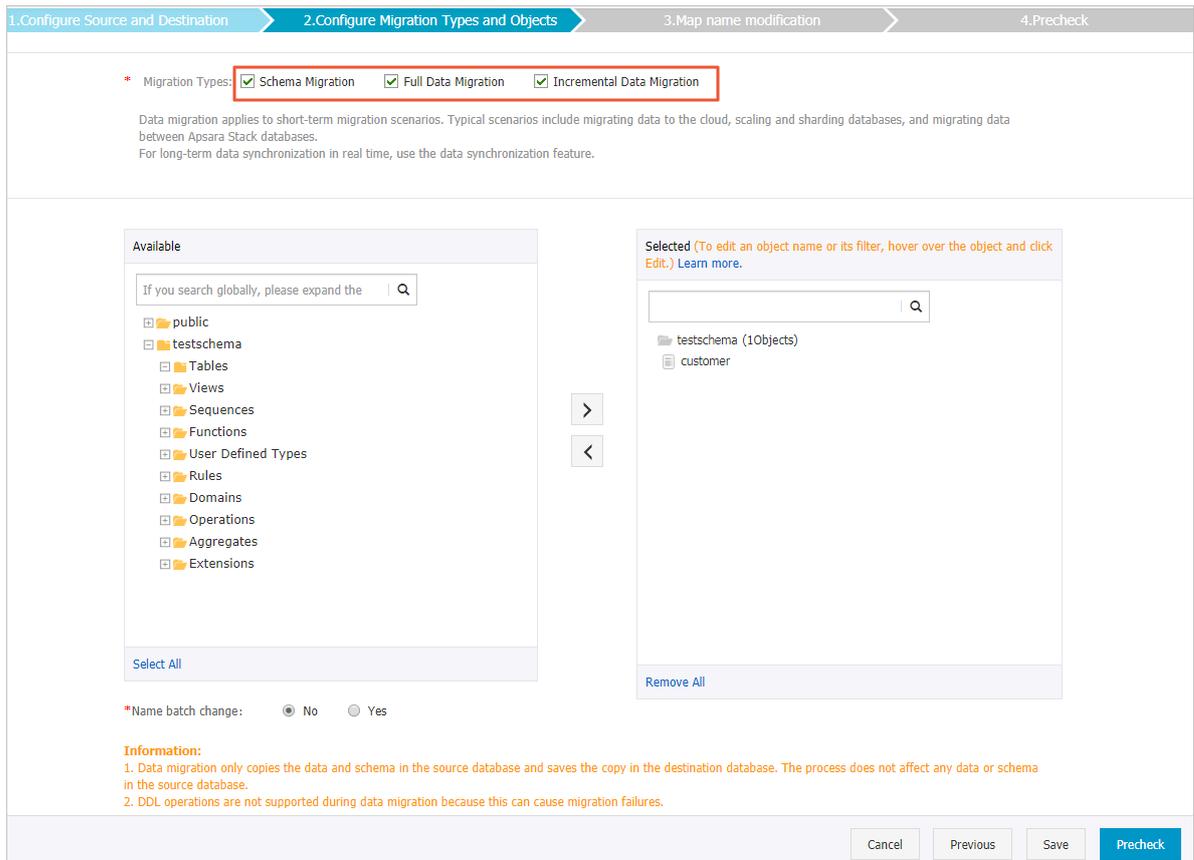
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



Note:

The CIDR blocks of DTS servers are automatically added to the whitelist of the destination 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.



Parameter	Description
Migration types	Select Schema Migration , Full Data Migration , and Incremental Data Migration .
Objects to be migrated	<p>Select objects from the Available section and click the  icon to add the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> 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 unchanged. You can change the name of an object in the destination ApsaraDB RDS for PostgreSQL instance by using the object name mapping feature provided by DTS. For more information about the authorization methods, 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **Data Transmission Service (Pay-As-You-Go) Service Terms**.

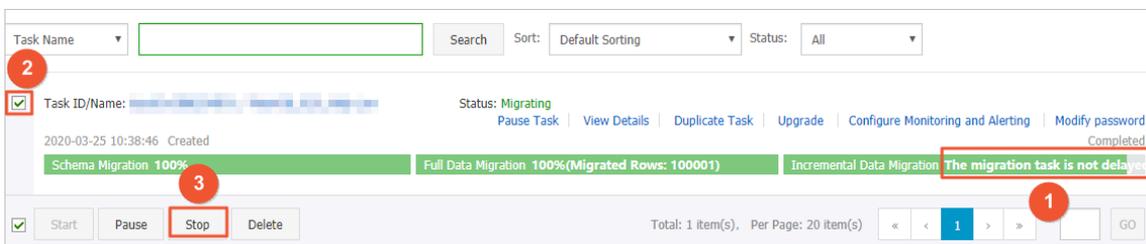
11. Click **Buy and Start** to start the migration task.



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.

- a. 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.
- 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 RDS for PostgreSQL instance.

7.5 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 PostgreSQL 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

- You cannot use DTS to migrate incremental data from an Amazon RDS for PostgreSQL database to an ApsaraDB RDS for PostgreSQL database. Before you start the data migration task, you must stop your business that runs on the Amazon RDS for PostgreSQL database. To ensure data consistency, do not write new data into the Amazon RDS for PostgreSQL database during data migration.
- 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.
- 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 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 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 #unique_51 .

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
ApsaraDB RDS for PostgreSQL database	The create and usage permissions on the objects to be migrated	The owner permission for schemas

Full data migration

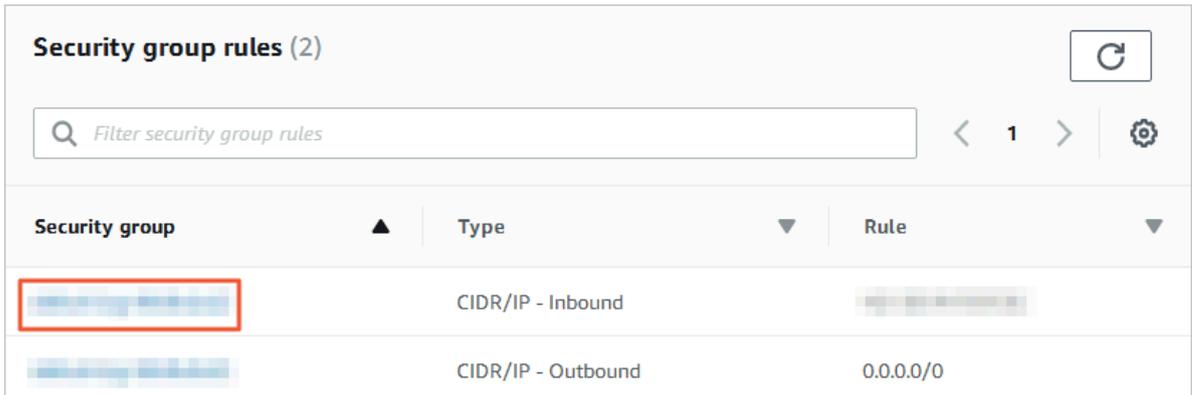
To avoid data migration failures caused by dependencies between objects, DTS migrates the schemas and data of the source PostgreSQL database in the following order:

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.

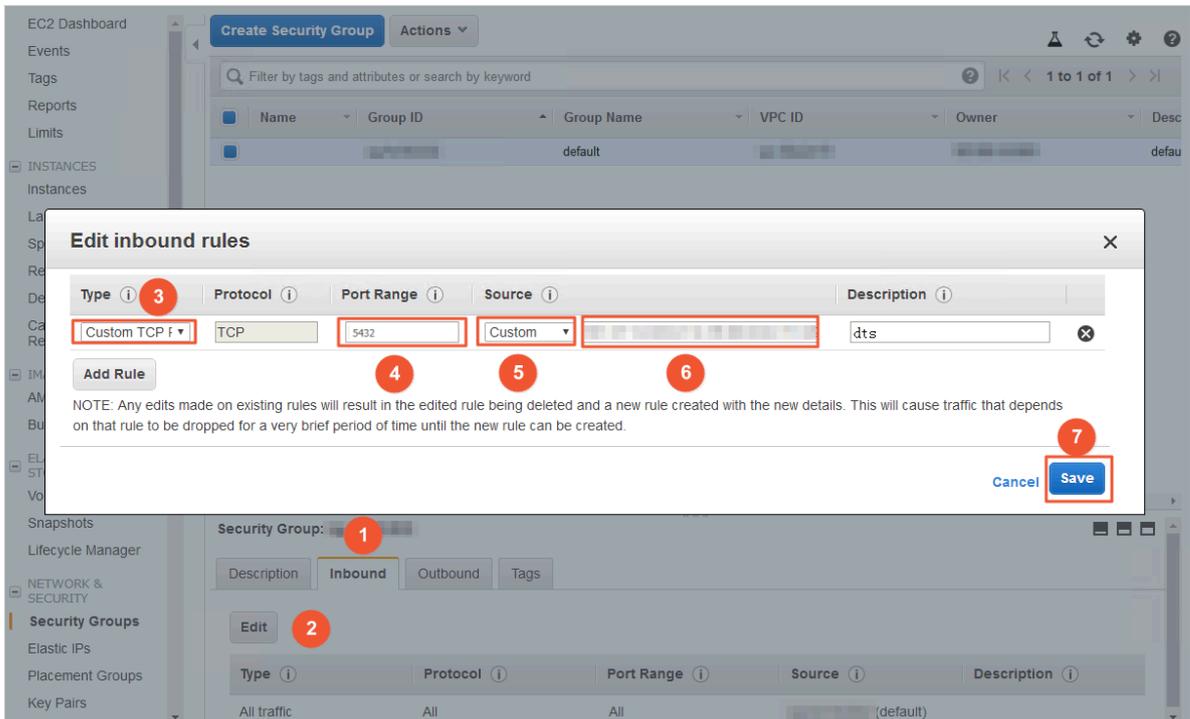


The screenshot shows the 'Security group rules (2)' interface. It includes a search bar with the placeholder 'Filter security group rules', a refresh button, and a table with the following data:

Security group	Type	Rule
[Redacted]	CIDR/IP - Inbound	[Redacted]
[Redacted]	CIDR/IP - Outbound	0.0.0.0/0

4. On the **Security Groups** page, click the Inbound tab in the Security Group section. On the Inbound tab, click Edit, 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 [#unique_62](#).



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.

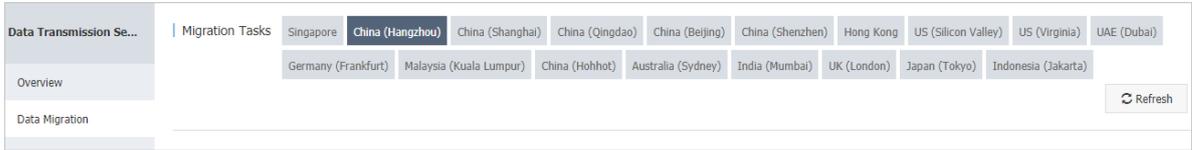
Preparation 2: Create a database and schema in the destination RDS instance

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

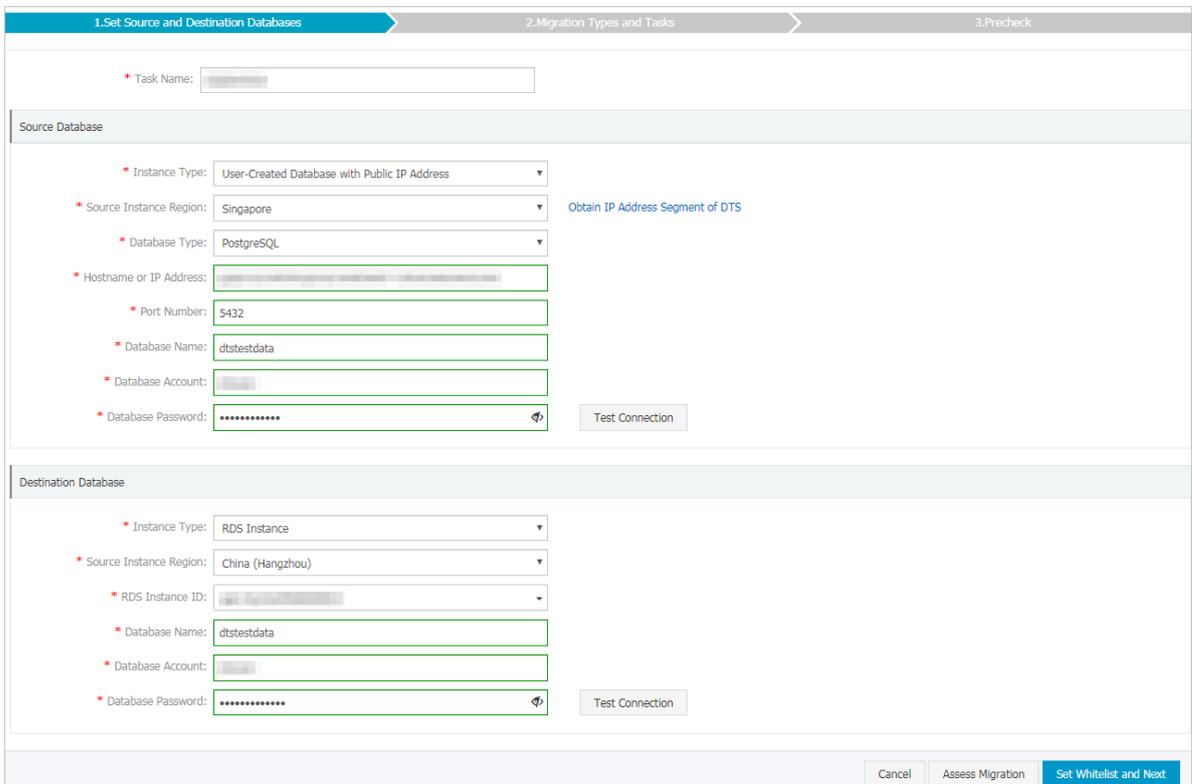
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.

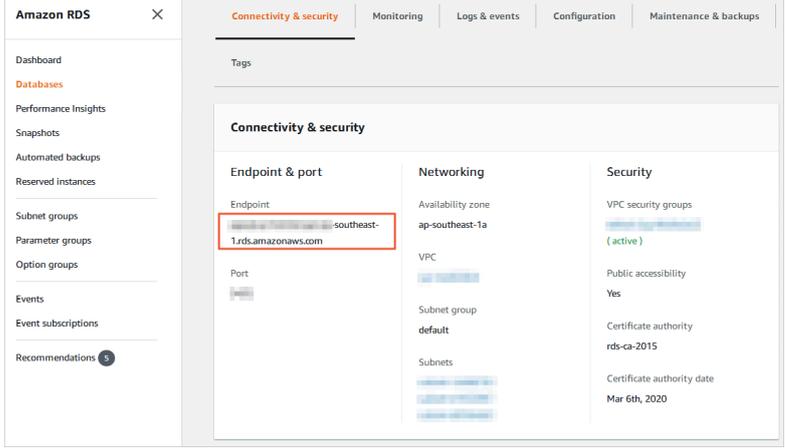


4. In the upper-right corner of the page, click **Create Migration Task**.

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.
Source Database	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
	Hostname or IP Address	<p>Enter the endpoint that is used to connect to the Amazon RDS for PostgreSQL database.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  Note: You can obtain the endpoint on the Basic information page of the source Amazon RDS for PostgreSQL instance. </div> 
	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.
	Database 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 .
	Database Password	<p>Enter the password for the database account.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  Note: 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 as prompted. </div>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.

Section	Parameter	Description
	RDS Instance ID	Select the ID of the destination ApsaraDB RDS for PostgreSQL instance.
	Database 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.  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 .
	Database 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 .
	Database Password	Enter the password for the database account.  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 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 RDS instance.

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

The screenshot shows the configuration interface for migration types and objects. At the top, there are four steps: 1. Configure Source and Destination, 2. Configure Migration Types and Objects (current), 3. Map name modification, and 4. Precheck.

Migration Types: Schema Migration, Full Data Migration, Incremental Data Migration. Below this, a note states: "During full data migration, data updates in the source database are not migrated to the destination instance. For data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration."

Available Objects: A tree view showing a hierarchy: public, testschema (expanded), Tables, Views, Sequences, Functions, User Defined Types, Rules, Domains, Operations, Aggregates, Extensions. A "Select All" button is at the bottom.

Selected Objects: A list showing "testschema (10Objects)" and "customer". A "Remove All" button is at the bottom.

Name batch change: No, Yes.

Information:
 1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.
 2. DDL operations are not supported during data migration because this can cause migration failures.

Buttons at the bottom: Cancel, Previous, Save, Precheck.

Parameter	Description
Migration types	<p>Select Schema Migration and Full Data Migration. In this data migration task, incremental data migration is not supported.</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p> Note: To ensure data consistency, do not write new data into the Amazon RDS for PostgreSQL database during data migration.</p> </div>

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> 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 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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 migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.



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

7.6 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.
- An ApsaraDB RDS for MySQL instance is created. For more information, see [Create an RDS for MySQL instance](#).
- 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.

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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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 data 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

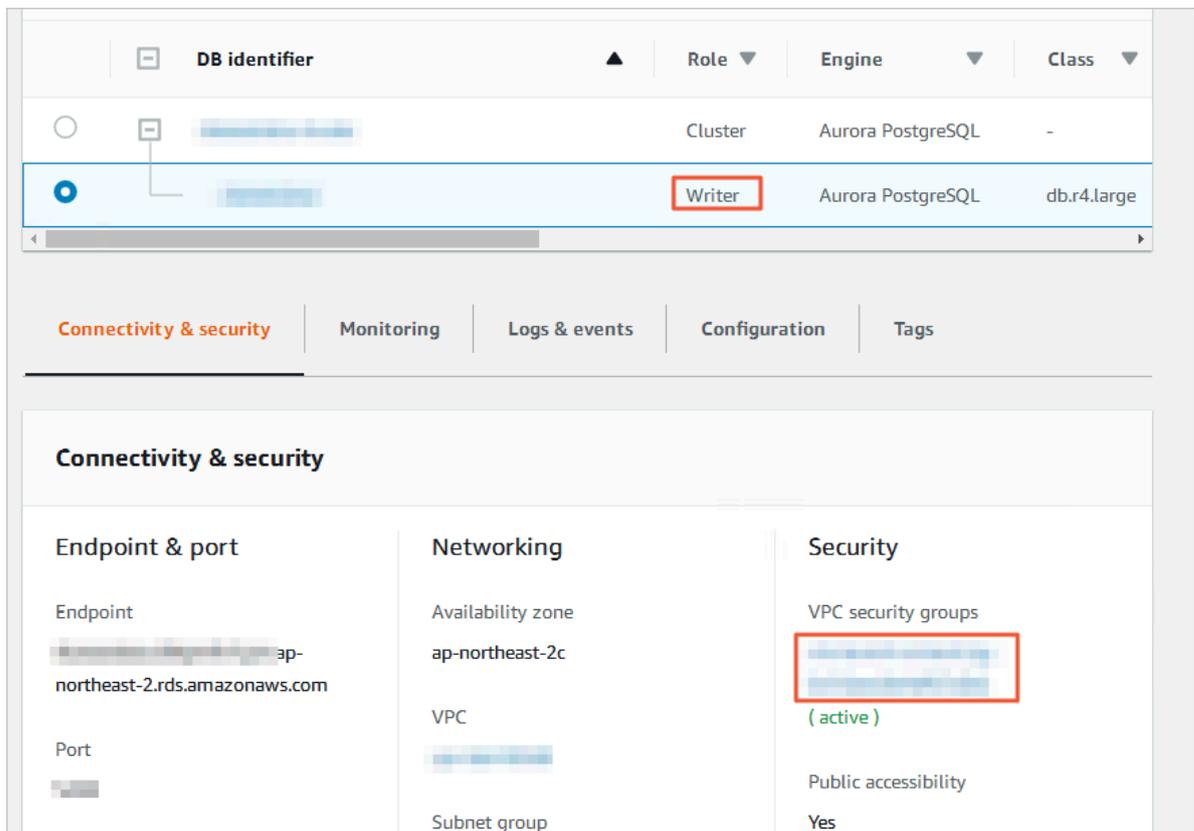
Database	Schema migration	Full data migration	Incremental data migration
ApsaraDB RDS for MySQL	The read/write permissions for the objects to be migrated	The read/write permissions for the objects to be migrated	The read/write permissions for the objects to be migrated

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

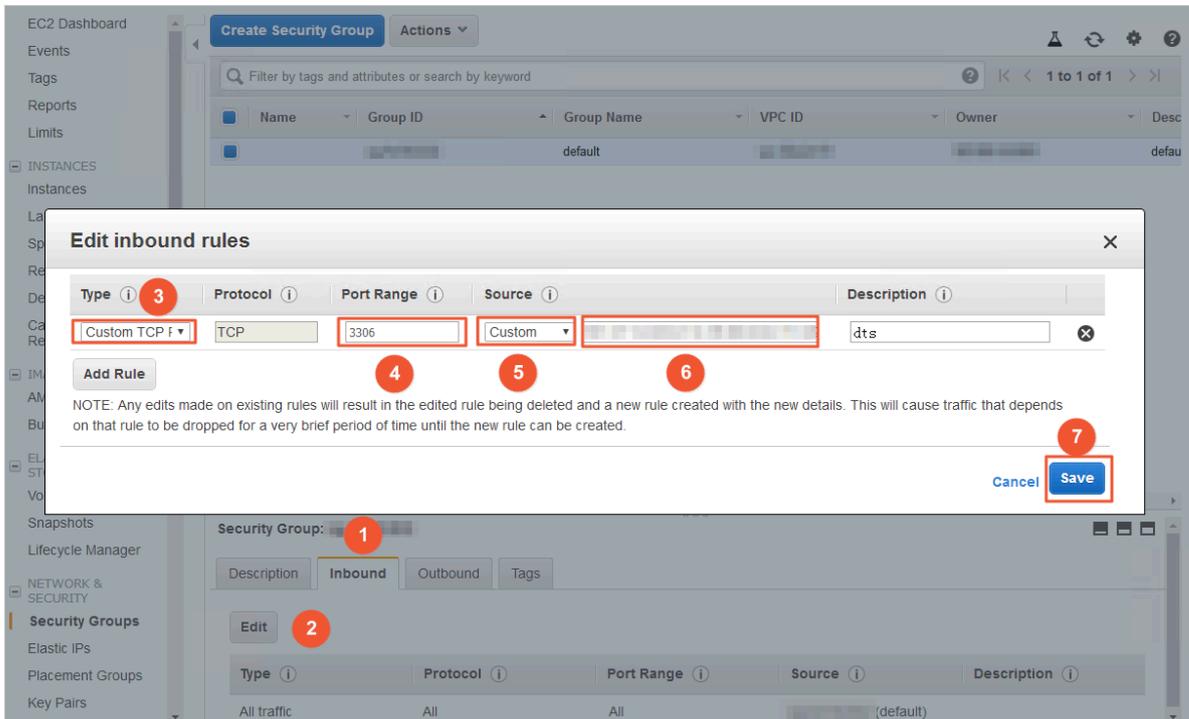
- [#unique_66](#) for an Amazon Aurora MySQL database
- [Create an account for an RDS for MySQL instance](#) and [Change the permissions of an account for an RDS for MySQL instance](#) for an ApsaraDB RDS for MySQL database

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.



5. 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 [#unique_62](#).



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.

6. Log on to the Amazon Aurora MySQL database and specify the number of hours to retain binary log files. Skip this step if you do not need to perform incremental data migration.

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



Note:

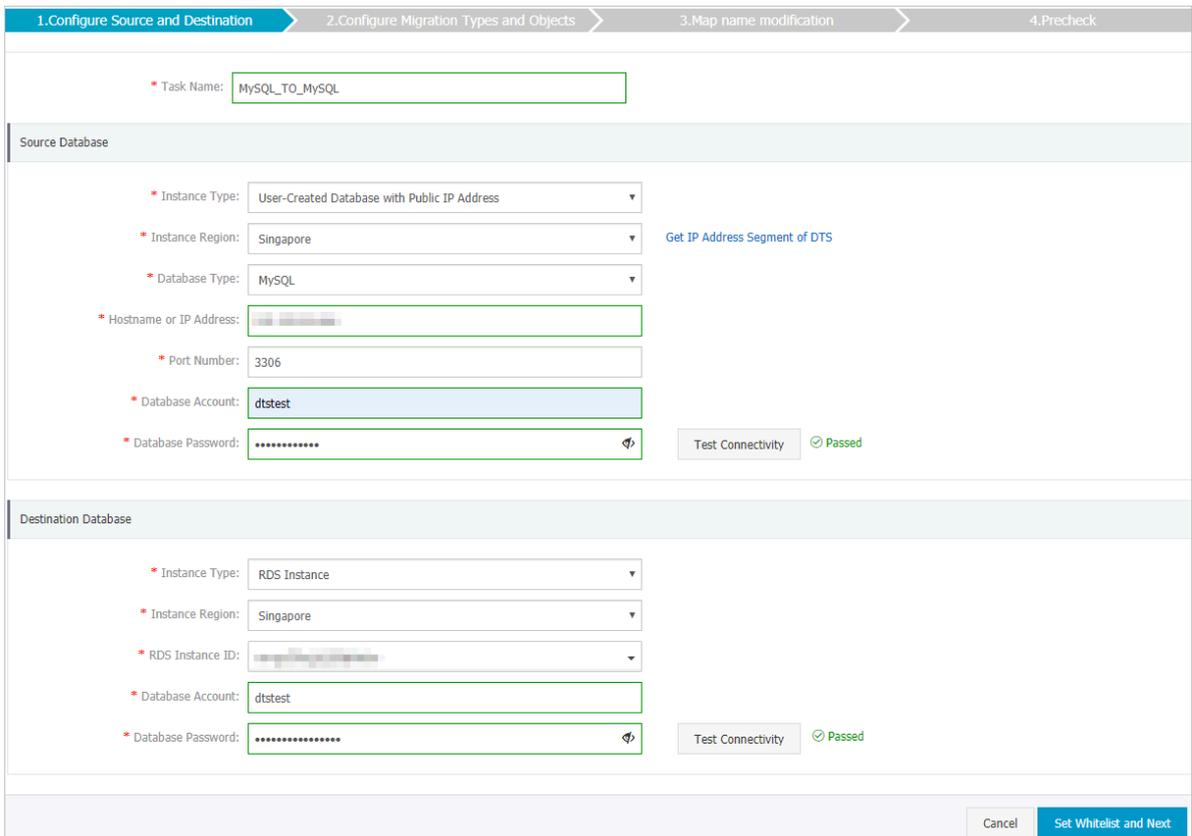
- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of Amazon Aurora MySQL must be enabled and the value of the binlog_format parameter must be set to row. If the version of MySQL is 5.6 or later, the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

Procedure

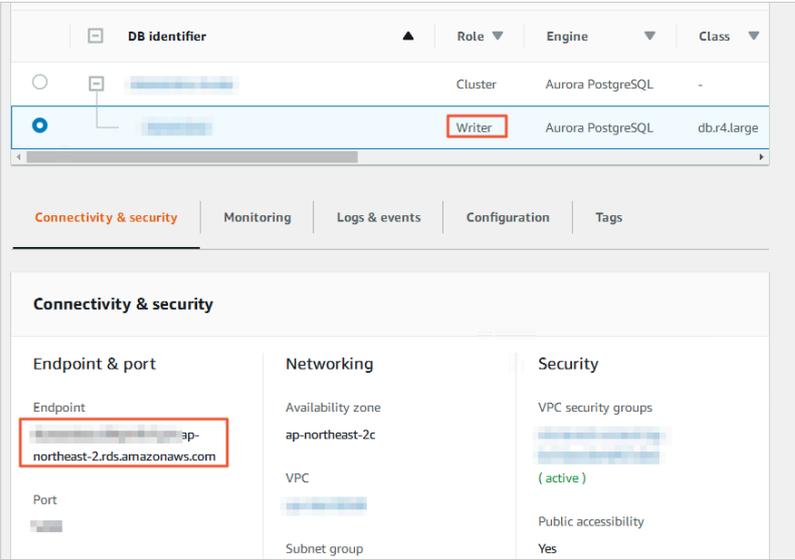
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.



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.

Section	Parameter	Description
Source Database	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 .
	Hostname or IP Address	Enter the endpoint that is used to connect to the Amazon Aurora MySQL database. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p> Note: You can obtain the endpoint on the Basic information page of the source Amazon Aurora MySQL instance.</p> </div> 
	Port Number	Enter the service port number of the Amazon Aurora MySQL database. The default port number is 3306 .
Database Account	Enter the account for the Amazon Aurora MySQL database. For more information about permissions required for the account, see Permissions required for database accounts .	

Section	Parameter	Description
	Database Password	<p>Enter the password for the database account.</p> <p> Note: After the source database parameters are specified, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are correct. If the source database parameters are correct, the Test Passed message is displayed, If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database parameters as prompted.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the ApsaraDB RDS for MySQL instance resides.
	RDS Instance ID	Select the ID of the ApsaraDB RDS for MySQL instance.
	Database 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 .
	Database Password	<p>Enter the password for the database account.</p> <p> Note: After the destination database parameters are specified, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are correct. If the destination database parameters are correct, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database parameters as prompted.</p>

Section	Parameter	Description
	Encryption	<p>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.</p> <p> Note: Encryption is available only in mainland China and Hong Kong(China).</p>

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 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.

The screenshot shows the '2. Configure Migration Types and Objects' step of the migration process. At the top, there are four progress indicators: '1. Configure Source and Destination', '2. Configure Migration Types and Objects' (highlighted), '3. Map name modification', and '4. Precheck'. Below the progress indicators, there are three checked options for migration types: 'Schema Migration', 'Full Data Migration', and 'Incremental Data Migration'. A note explains that data migration is for short-term scenarios, while data synchronization is for long-term real-time scenarios. The main area is split into two panels: 'Available' on the left and 'Selected' on the right. The 'Available' panel shows a tree view of database objects including 'data123', 'mysqltest' (with sub-items like Tables, Views, Functions, Procedures), and 'sys'. The 'Selected' panel shows 'mysqltest (2 Objects)' with sub-items 'customer' and 'vipinfo'. At the bottom, there are radio buttons for 'Name batch change' (set to 'No') and an 'Information' section with two warning points. Navigation buttons 'Cancel', 'Previous', 'Save', and 'Precheck' are at the bottom right.

Parameter	Description
Migration types	<ul style="list-style-type: none"> 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. <p> 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.</p>

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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.

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

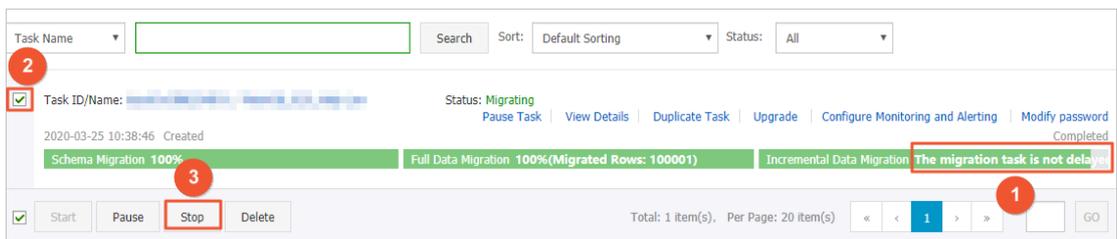
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When 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 MySQL instance.

7.7 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 Amazon Aurora MySQL is set to **Yes**. This ensures that DTS can access Amazon Aurora MySQL 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 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%.
- 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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

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.



Note:

- 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

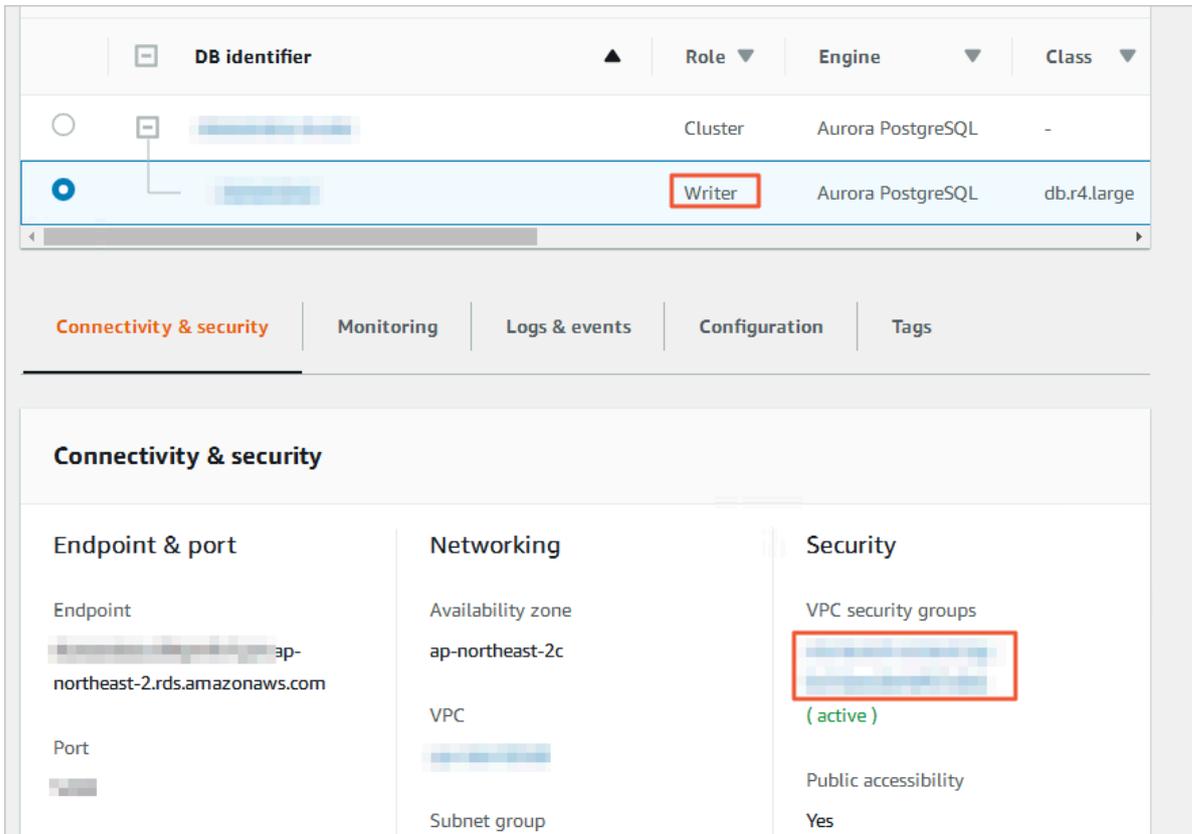
For more information about how to create and authorize a database account, see the following topics:

- [#unique_66](#) for an Amazon Aurora MySQL database
- [Create database accounts](#) for a PolarDB for MySQL database

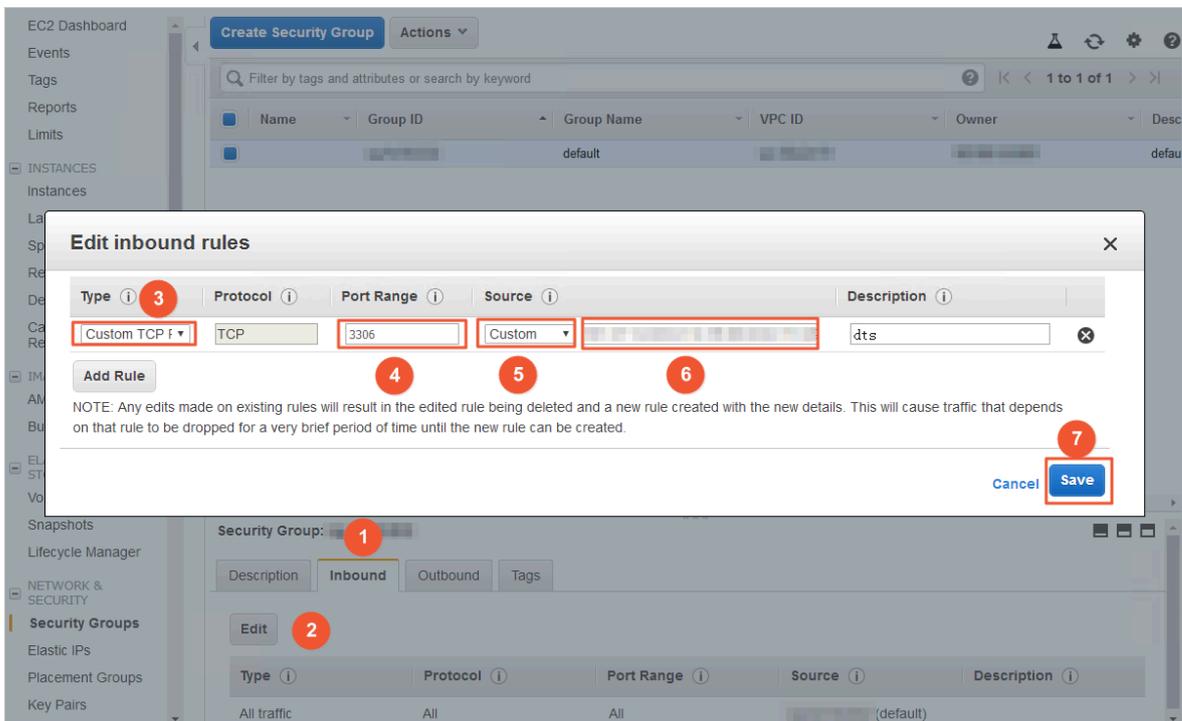
Preparations before data migration

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.



5. 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 [#unique_62](#).



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.

6. Log on to the Amazon Aurora MySQL database and specify the number of hours to retain binary log files. Skip this step if you do not need to perform incremental data migration.

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



Note:

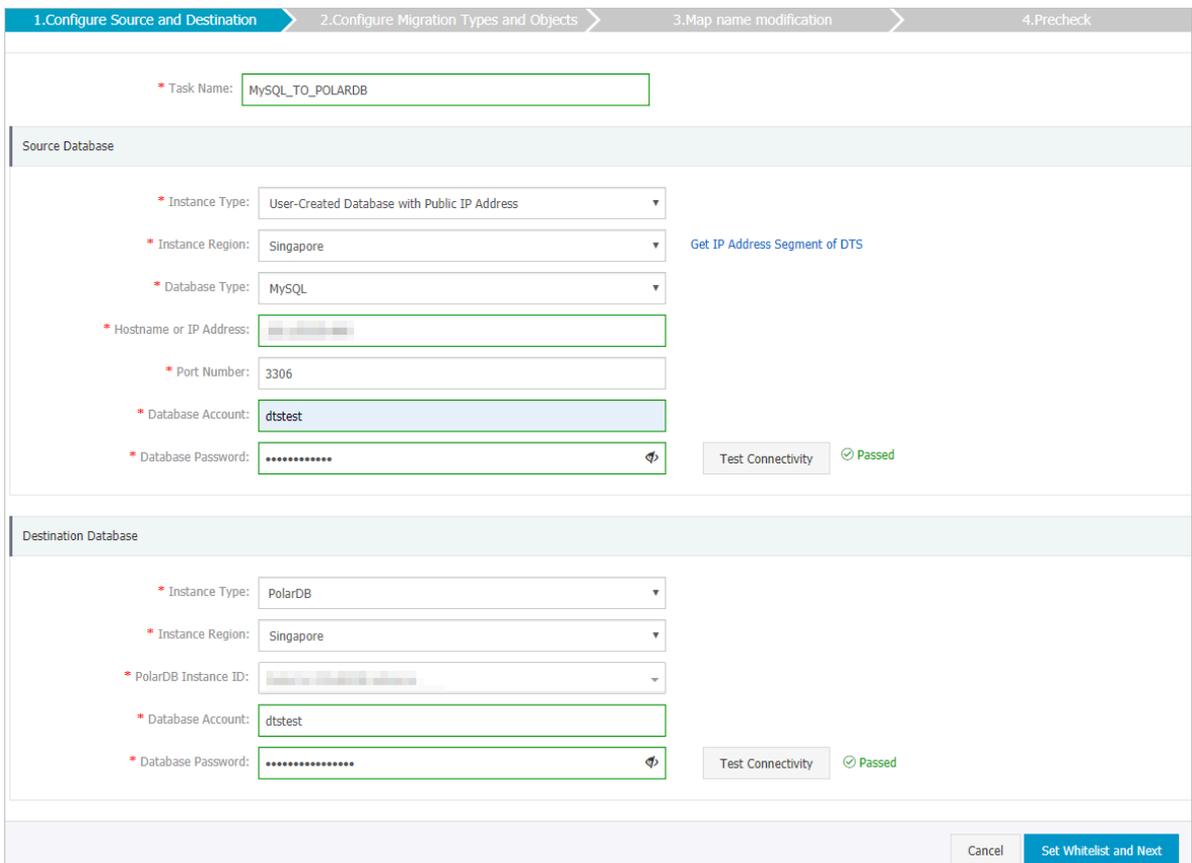
- The preceding command sets the retention period of binary log files to 24 hours. The maximum value is 168 hours (7 days).
- The binary logging feature of Amazon Aurora MySQL must be enabled and the value of the binlog_format parameter must be set to row. If the version of MySQL is 5.6 or later, the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

Procedure

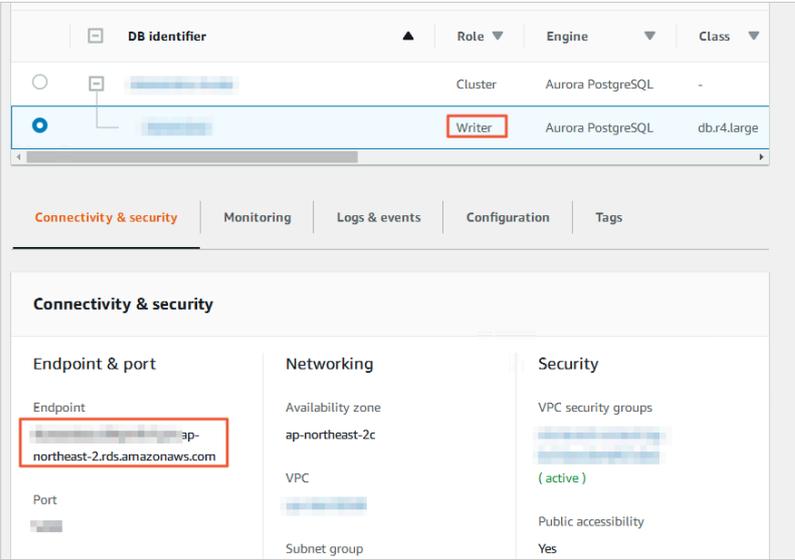
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.



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.

Section	Parameter	Description
Source Database	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 .
	Hostname or IP Address	<p>Enter the endpoint that is used to connect to the Amazon Aurora MySQL database.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;">  Note: You can obtain the endpoint on the Basic information page of the source Amazon Aurora MySQL instance. </div> 
	Port Number	Enter the service port number of the Amazon Aurora MySQL database. The default port number is 3306 .
Database Account	Enter the account for the Amazon Aurora MySQL database. For more information about permissions required for the account, see Permissions required for database accounts .	

Section	Parameter	Description
	Database Password	<p>Enter the password for the database account.</p> <p> Note: 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 as prompted.</p>
Destination Database	Instance Type	Select PolarDB .
	Instance Region	Select the region where the ApsaraDB RDS for MySQL instance resides.
	PolarDB instance ID	Select the ID of the PolarDB for MySQL instance.
	Database Account	Enter the database account of the PolarDB for MySQL instance. For more information about permissions required for the account, see Permissions required for database accounts .
	Database Password	<p>Enter the password for the database account.</p> <p> 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 as prompted.</p>

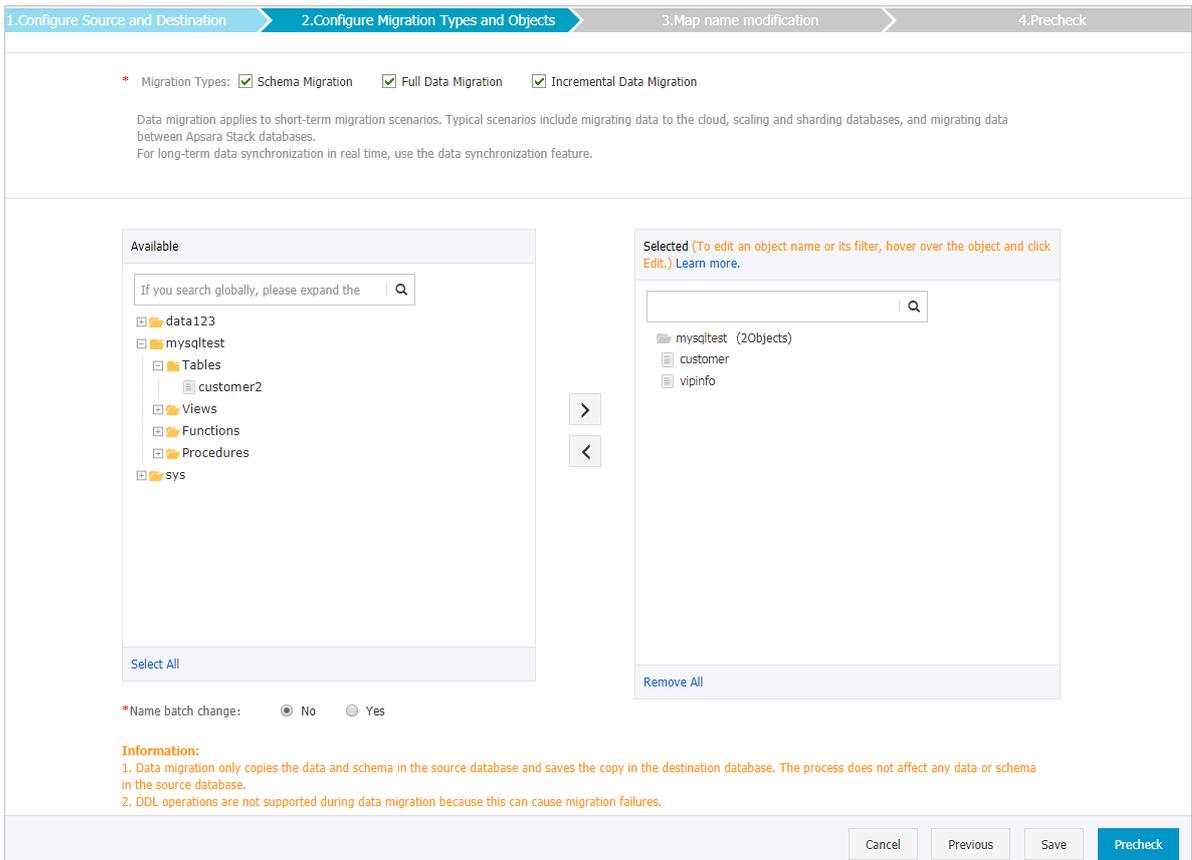
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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.



Parameter	Description
Migration types	<ul style="list-style-type: none"> 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. <p> 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.</p>

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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.

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

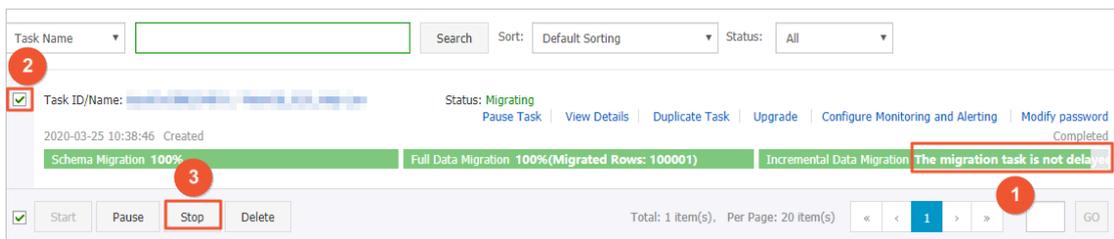
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When 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 PolarDB for MySQL instance.

7.8 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 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 Aurora PostgreSQL instance.

Notes

- You cannot use DTS to migrate incremental data from an Amazon Aurora PostgreSQL database to an ApsaraDB RDS for PostgreSQL database. Before you start the data migration task, you must stop your business that runs on the Amazon Aurora PostgreSQL database. To ensure data consistency, do not write new data into the Amazon Aurora PostgreSQL database during data migration.
- 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.
- 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 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 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 #unique_51 .

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 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

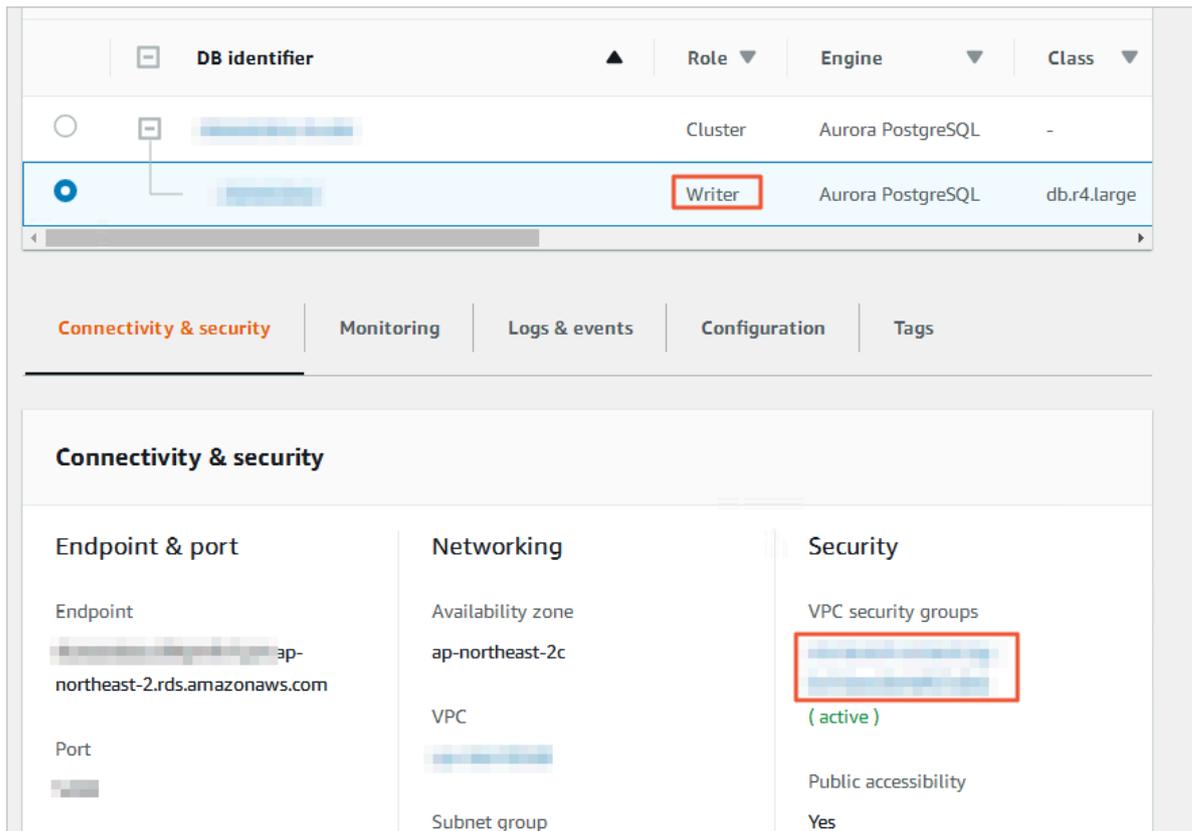
Full data migration

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.

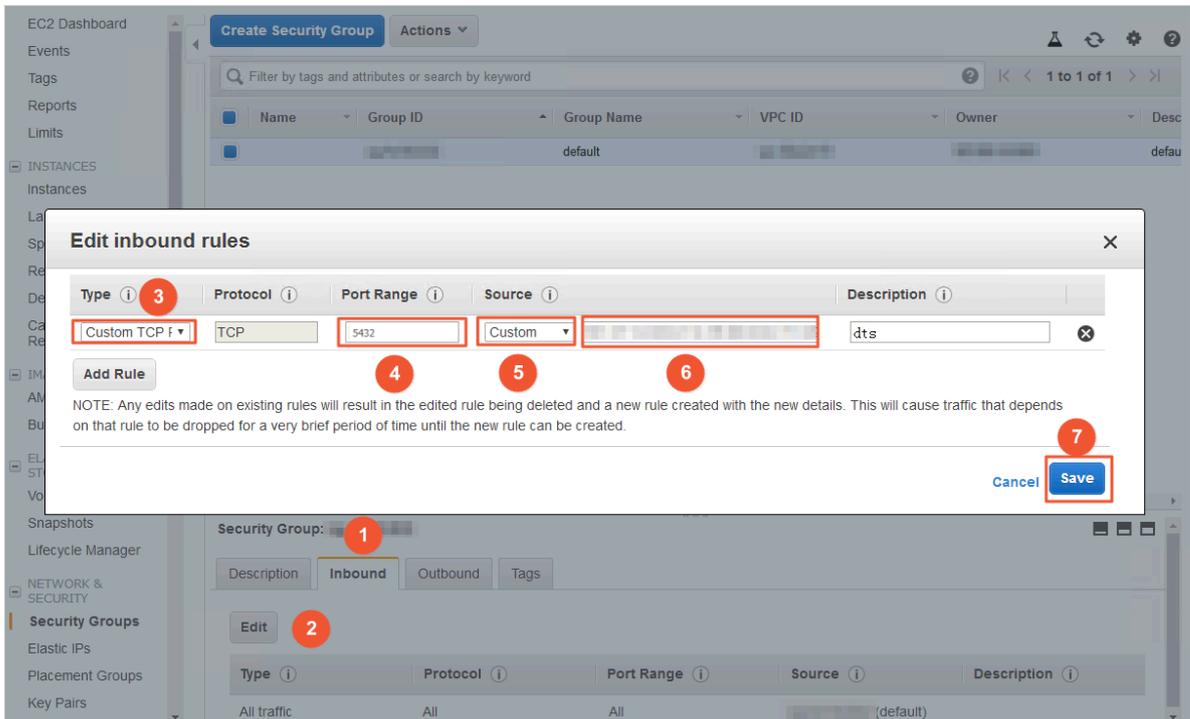
Preparation 1: Modify Amazon inbound rules

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



5. 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 [#unique_62](#).



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.

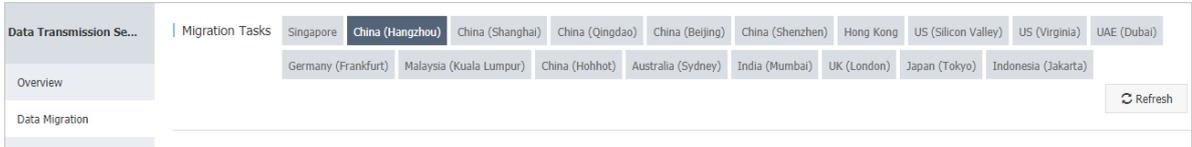
Preparation 2: Create a database and schema in the destination RDS instance

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

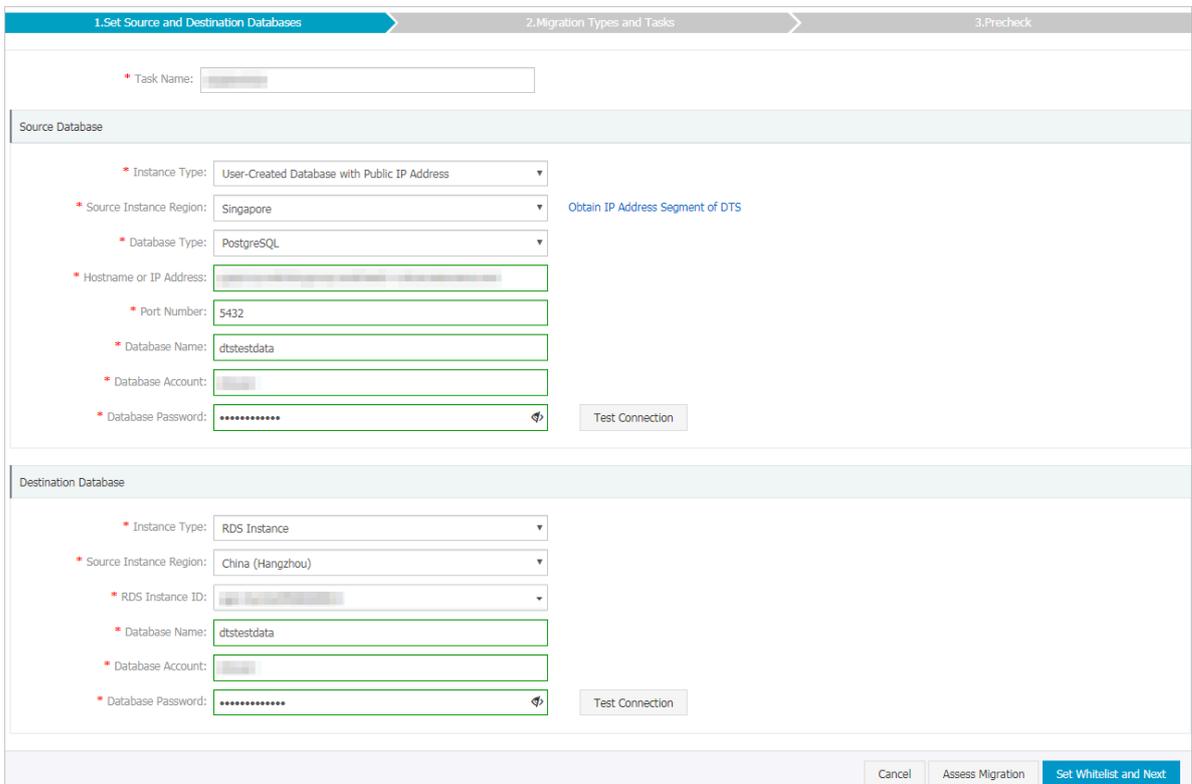
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of **Migration Tasks** the page, select the region where the destination cluster resides.

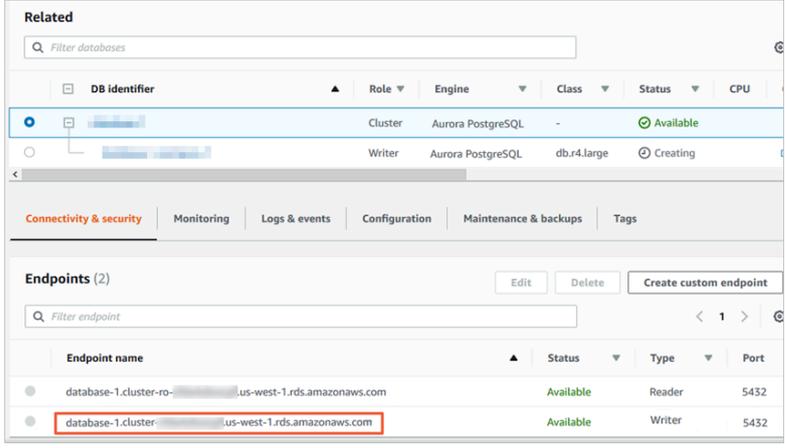


4. In the upper-right corner of the page, click **Create Migration Task**.

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.
Source Database	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
	Hostname or IP Address	<p>Enter the endpoint that is used to connect to the Amazon Aurora PostgreSQL database.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  Note: You can obtain the endpoint on the Basic information page of the source Amazon Aurora PostgreSQL instance. </div> 
	Port Number	Enter the service port number of the Amazon Aurora PostgreSQL database. The default port number is 5432 .
	Database Name	Enter the name of the source Amazon Aurora PostgreSQL database.
	Database 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 .
	Database Password	<p>Enter the password for the database account.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  Note: 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 as prompted. </div>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination ApsaraDB RDS for PostgreSQL instance resides.

Section	Parameter	Description
	RDS Instance ID	Select the ID of the destination ApsaraDB RDS for PostgreSQL instance.
	Database Name	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.  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 .
	Database 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 .
	Database Password	Enter the password for the database account.  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 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 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.

1. Configure Source and Destination
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Migration Types: Schema Migration Full Data Migration Incremental Data Migration

During full data migration, data updates in the source database are not migrated to the destination instance. For data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration.

Available

If you search globally, please expand the

- public
- testschema
 - Tables
 - Views
 - Sequences
 - Functions
 - User Defined Types
 - Rules
 - Domains
 - Operations
 - Aggregates
 - Extensions

Select All

Selected (To edit an object name or its filter, hover over the object and click Edit.) [Learn more.](#)

- testschema (10Objects)
 - customer

Remove All

*Name batch change: No Yes

Information:
 1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.
 2. DDL operations are not supported during data migration because this can cause migration failures.

Cancel Previous Save Precheck

Parameter	Description
Migration types	Select Schema Migration and Full Data Migration . In this data migration task, incremental data migration is not supported.
	<div style="display: flex; align-items: center; margin-bottom: 5px;"> <p>Note:</p> </div> <p>To ensure data consistency, do not write new data into the Amazon Aurora PostgreSQL database during data migration.</p>

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> You can select databases, tables, or columns as the 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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 migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.



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

7.9 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 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%.

- Amazon RDS for SQL Server does not support incremental data migration because the sysadmin server-level role is not available in Amazon RDS.
- 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.
- 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 #unique_51 .

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

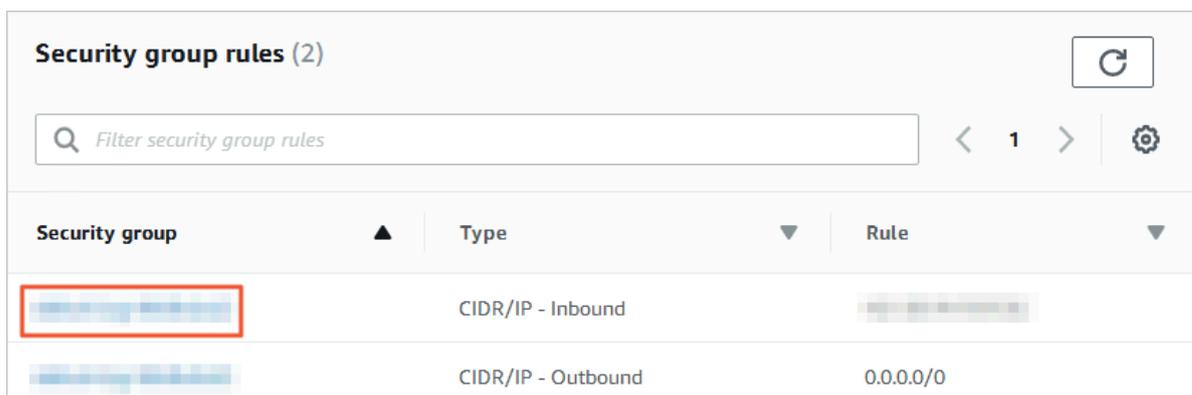
Full data migration

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:

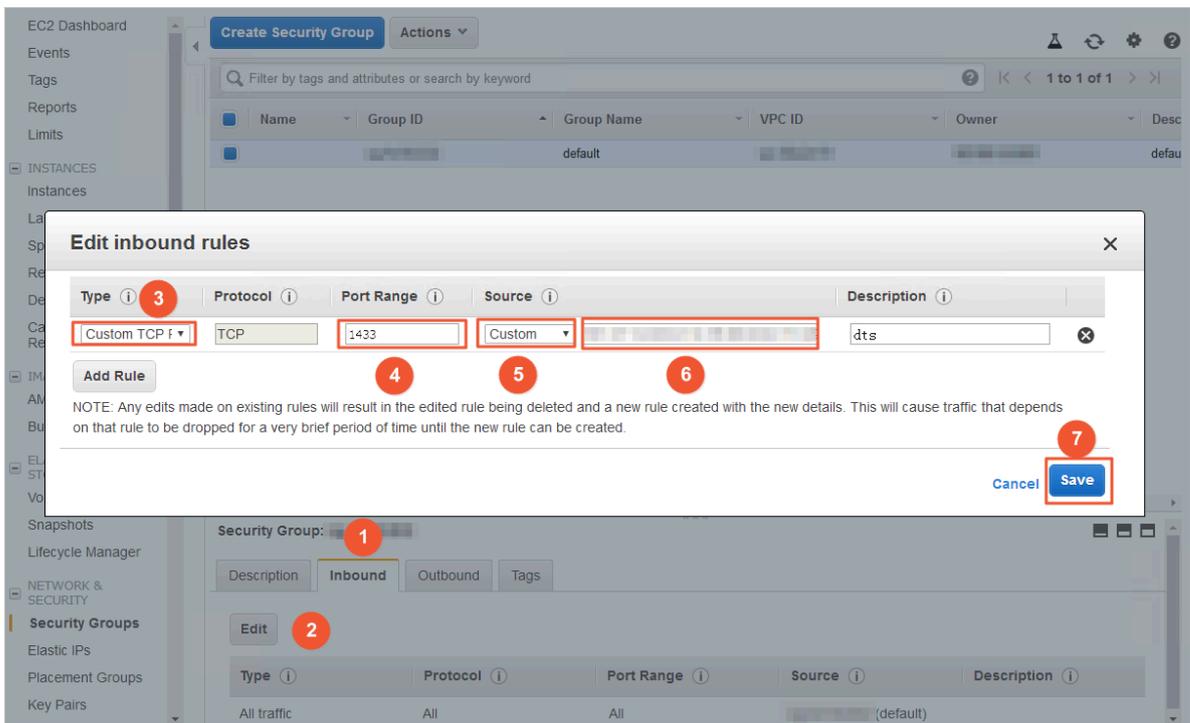
1. Migrate the schemas of tables, views, synonyms, user-defined types, rules, defaults, and plan guides.
2. Perform full data migration.
3. Migrate the schemas of SQL stored procedures, SQL functions, triggers, and foreign keys.

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 [#unique_62](#).



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 **Migration Tasks** the page, select the region where the destination cluster resides.

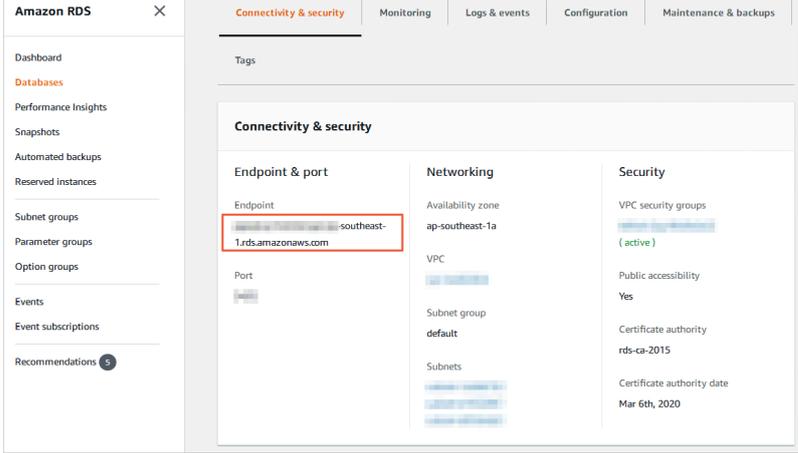


4. In the upper-right corner of the page, click **Create Migration Task**.

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

The screenshot shows the configuration interface for a data migration task, divided into four steps: 1. Configure Source and Destination, 2. Configure Migration Types and Objects, 3. Map name modification, and 4. Precheck. The current step is 1. The 'Task Name' is 'SQL_Server_TO_RDS'. The 'Source Database' section includes fields for Instance Type (User-Created Database with Public IP Address), Instance Region (Singapore), Database Type (SQLServer), Hostname or IP Address, Port Number (1433), Database Account (dtstest), and Database Password. A 'Test Connectivity' button shows a 'Passed' status. The 'Destination Database' section includes fields for Instance Type (RDS Instance), Instance Region (Singapore), RDS Instance ID, Database Account (dtstest), and Database Password. A 'Test Connectivity' button also shows a 'Passed' status. At the bottom right, there are 'Cancel' and 'Set Whitelist and Next' buttons.

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 Database	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 SQL Server .

Section	Parameter	Description
	Hostname or IP Address	<p>Enter the endpoint that is used to connect to the Amazon RDS for SQL Server database.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Note: You can obtain the endpoint on the Basic information page of the source Amazon RDS for SQL Server instance.</p> </div> 
	Port Number	Enter the service port number of the Amazon RDS for SQL Server database. The default port number is 1433 .
	Database 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 .
	Database Password	<p>Enter the password for the database account.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>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.</p> </div>
Destination Database	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.

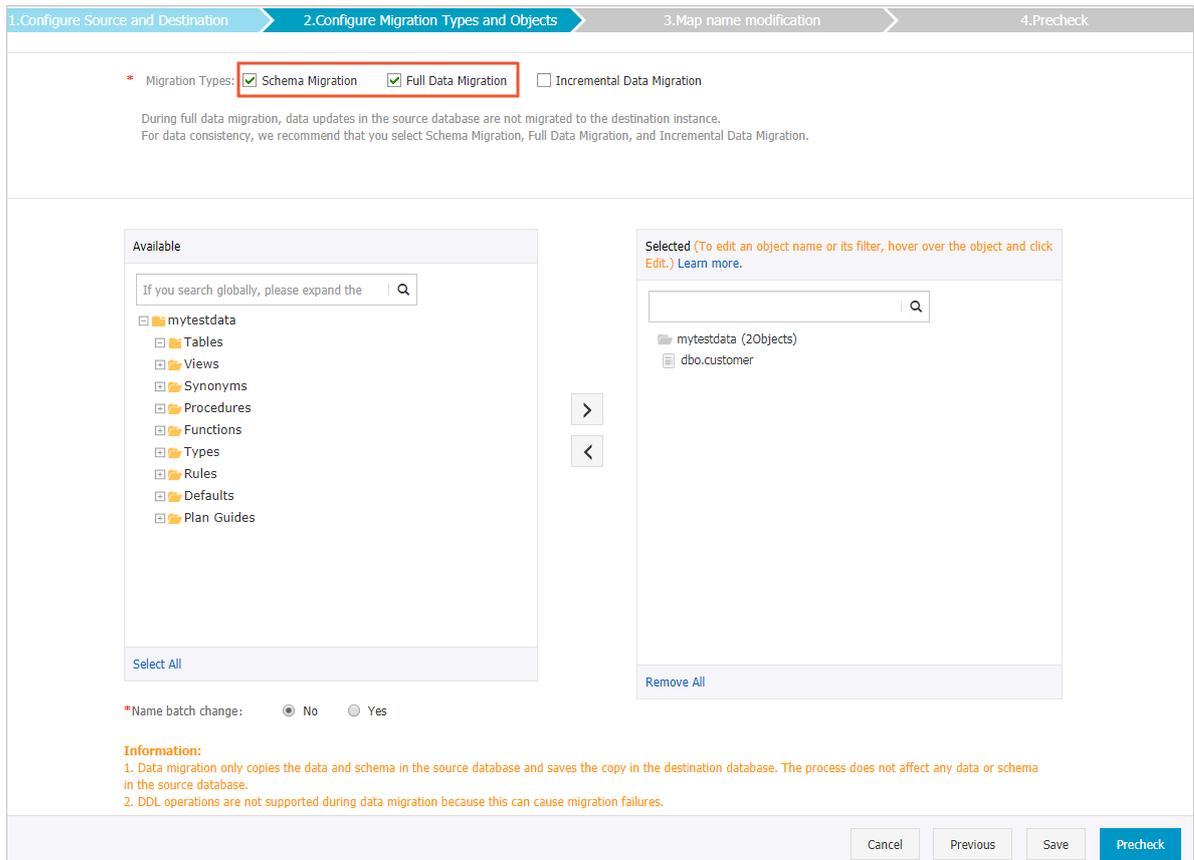
Section	Parameter	Description
	Database Account	Enter the database account of the ApsaraDB RDS for SQL Server instance. For more information about permissions required for the account, see Permissions required for database accounts .
	Database Password	Enter the password for the database account.  Note: After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the 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 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.



Parameter	Description
Migration types	<p>Select Schema Migration and Full Data Migration.</p> <p> Note:</p> <ul style="list-style-type: none"> Amazon RDS for SQL Server does not support incremental data migration because the sysadmin server-level role is not available in Amazon RDS. To ensure data consistency, do not write new data into Amazon RDS for SQL Server during full data migration.

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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 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. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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 migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.



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

7.10 Migrate a MongoDB Atlas database to Alibaba Cloud through DTS

This topic describes how to incrementally migrate a MongoDB Atlas database to Alibaba Cloud through Data Transmission Service (DTS). DTS supports the full migration and incremental migration methods. You can use both methods to migrate the MongoDB Atlas database without business interruption.

Prerequisites

The storage space of ApsaraDB for MongoDB instances must be greater than that of the MongoDB Atlas database.

Precautions

- We recommend that you migrate your data during off-peak hours to prevent negative impacts on your business.
- Data in the admin and local database is not migrated even if the admin database is selected.
- The config database is an internal database. Do not migrate this database unless otherwise essential to your business.
- For more information about the versions and storage engines supported by ApsaraDB for MongoDB instances, see [#unique_79](#). If you want to migrate the database between different versions or engines, confirm compatibility of these versions and engines first.

Billing

Migration type	Configuration fee	Public traffic fee
Full data migration	Free of charge.	You are charged for migrating data from the Internet to Alibaba Cloud. For more information, see #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

Migration types

Migration type	Description
Full data migration	<p>All data of the migration objects is migrated from the source instance to the destination instance.</p> <p> Note: Data migration is supported at the database, collection, and index levels.</p>
Incremental data migration	<p>Updated data of the migration objects is synchronized from the source instance to the destination instance.</p> <p> Note:</p> <ul style="list-style-type: none"> The create and delete operations on databases, collections, and indexes can be synchronized. The create, delete, and update operations on documents can be synchronized.

Account permissions for the database

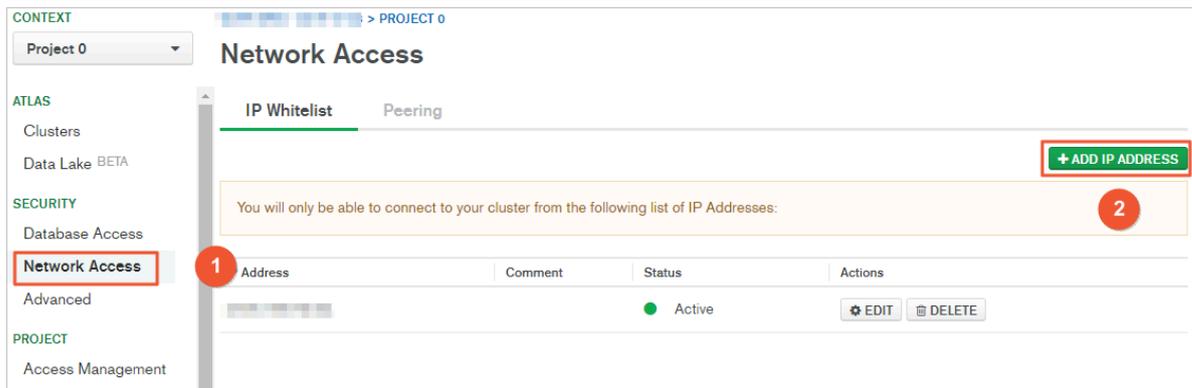
Database	Full data migration	Incremental data migration
MongoDB Atlas database	Permissions to read the database to be migrated and perform the listDatabases operation.	<ul style="list-style-type: none"> Permissions to read the database to be migrated, admin database, and local database. Permissions to perform the listDatabases operation.
ApsaraDB for MongoDB instance	Permissions to read and write the destination database.	Permissions to read and write the destination database.

How to create a database account and grant permissions to the account:

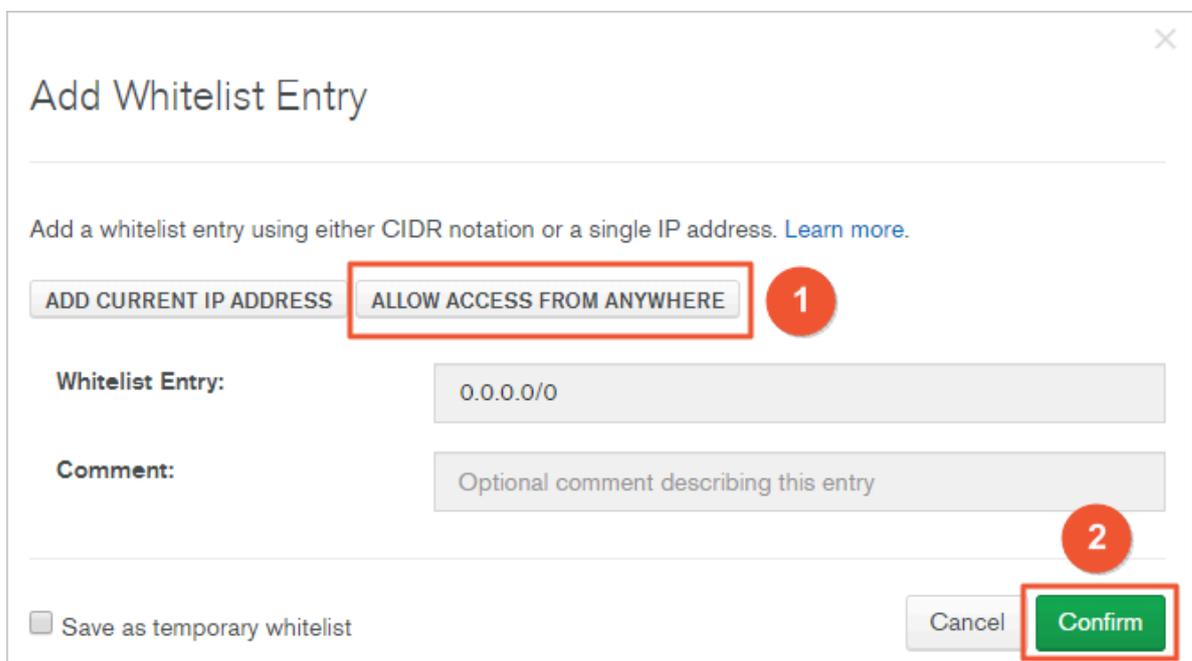
- For the MongoDB Atlas database, see [Create User in MongoDB](#).
- For the ApsaraDB for MongoDB instance, see [Manage MongoDB users through DMS](#).

Preparations

1. Log on to the MongoDB Atlas console.
2. In the left-side navigation pane, click **Network Access**. Click **ADD IP ADDRESS** on the page.



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



Note:

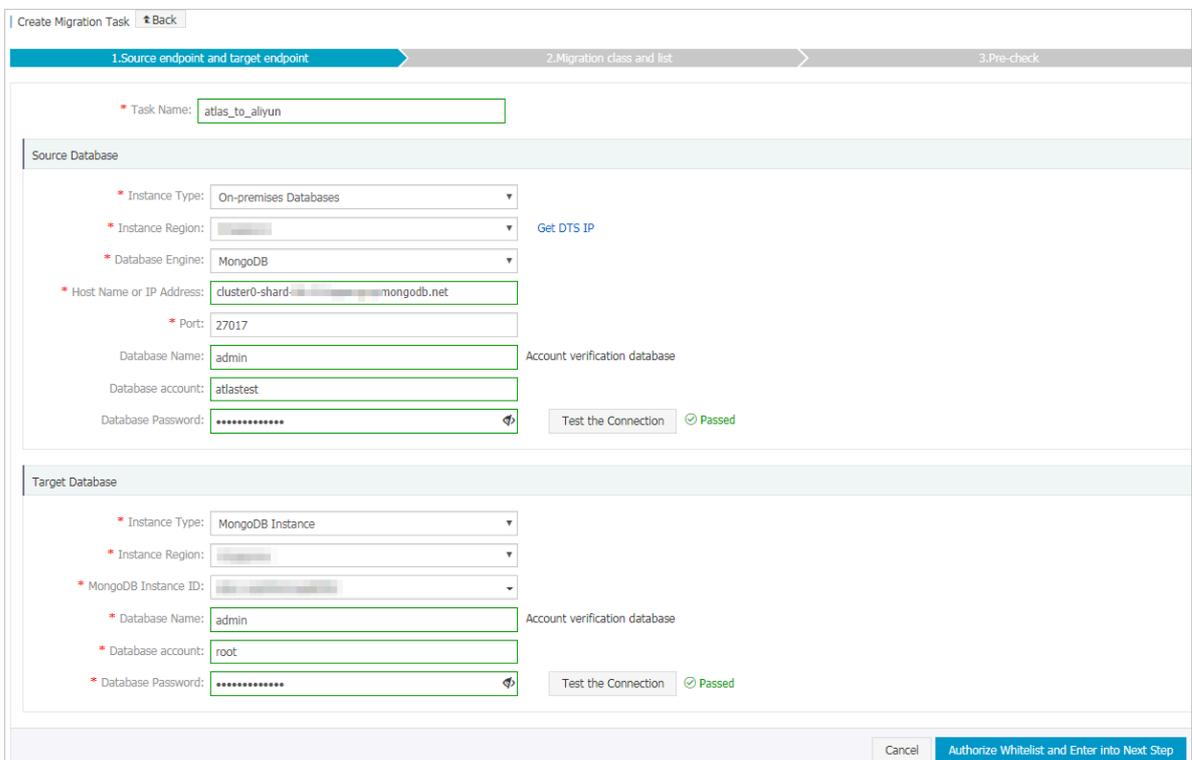
This step allows all IP addresses to access the MongoDB Atlas database. Delete this rule after the migration.

Procedure

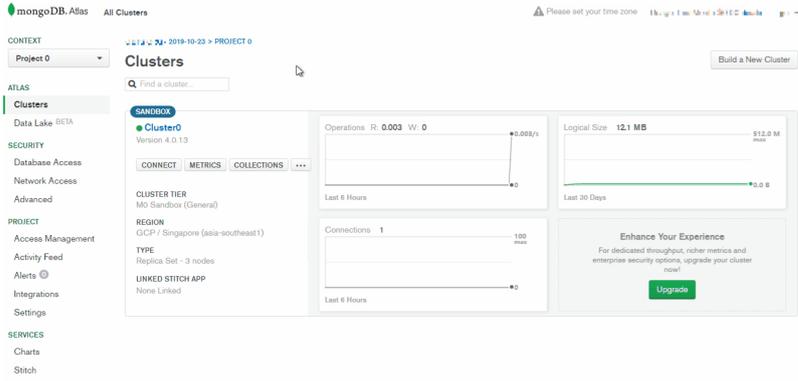
1. Log on to the [Data Transmission Service console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of the **Migration Tasks** page, select the region where the ApsaraDB for MongoDB instance is located.



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

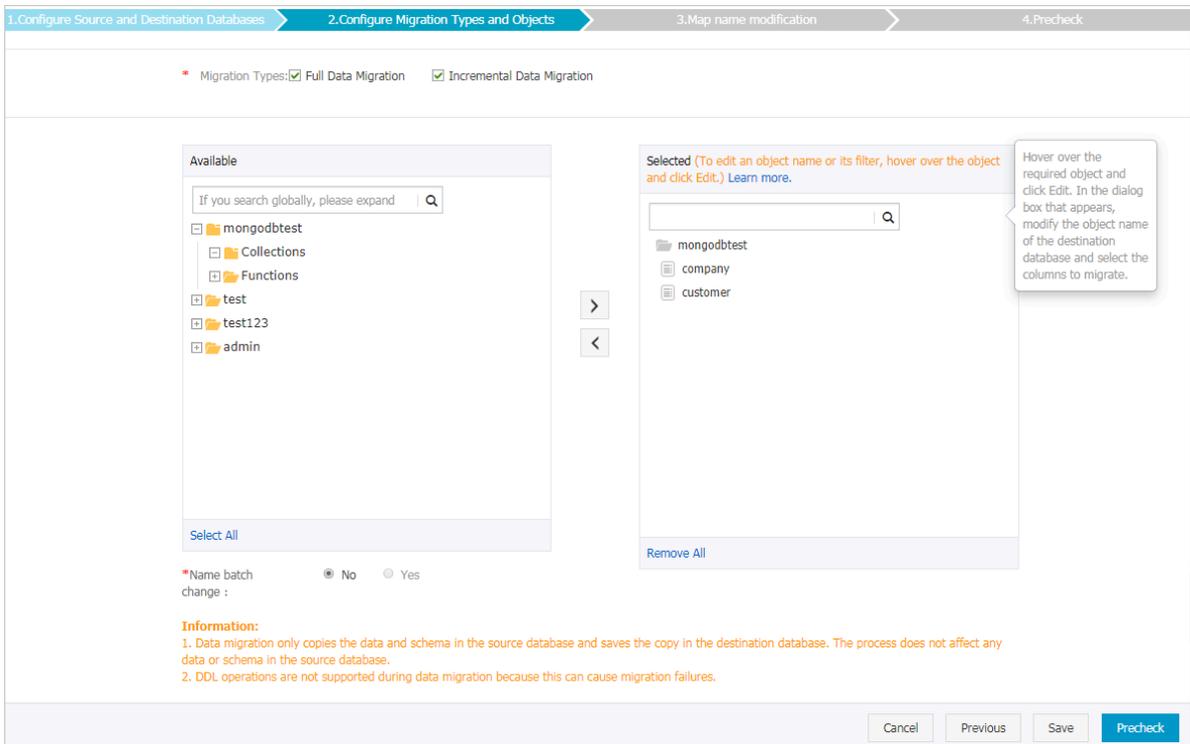


Section	Parameter	Description
Task Name	N/A	DTS automatically generates a task name. We recommend that you use an identifiable name to simplify future management. The name is not necessarily unique.
Source Database	Instance Type	Select User-Created Database with Public IP Address .
	Instance Region	When Instance Type is set to User-Created Database with Public IP Address , you do not need to set Instance Region .
	Database Type	Select MongoDB .

Section	Parameter	Description
	Hostname or IP Address	<p>Enter the endpoint of the PRIMARY node in the MongoDB Atlas database.</p> <p>You can obtain the endpoint in the MongoDB Atlas console , as shown in the following figure.</p> 
	Port Number	Enter the service port of the MongoDB Atlas database. The default port is 27017 .
	Database Name	Enter the name of the database used for account authentication.
	Database Account	Enter the account of the MongoDB Atlas database. For more information about permission requirements, see Account permissions for the database .
	Database Password	<p>Enter the password of the database account.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <p>After the source database configuration is complete, you can click Test Connectivity after the Database Password field to check whether the source database configuration is correct. If the source database configuration is correct, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database configuration as prompted.</p> </div>
	Encryption	Select SSL-encrypted .
Destination Database	Instance Type	Select MongoDB Instance .
	Instance Region	Select the region of the ApsaraDB for MongoDB instance.

Section	Parameter	Description
	MongoDB Instance ID	Select the ID of the ApsaraDB for MongoDB instance.
	Database Name	Enter the name of the database used for account authentication.  Note: If you use the root account, the database name is admin.
	Database Account	Enter the database account of the ApsaraDB for MongoDB instance. For more information about permission requirements, see Account permissions for the database .
	Database Password	Enter the password of the database account.  Note: After the destination database configuration is complete, you can click Test Connectivity after the Database Password field to check whether the destination database configuration is correct. If the destination database configuration is correct, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database configuration as prompted.

6. After you configure the parameters, click **Set Whitelist and Next** in the lower-right corner.



 **Note:**
Then the IP address of the DTS server is automatically added to the whitelist of the ApsaraDB for MongoDB instance. This ensures that the DTS server can connect to the ApsaraDB for MongoDB instance.

7. Configure the migration object and migration type.

Item	Description
Migration type	<ul style="list-style-type: none"> If you want to migrate all data, select Full Data Migration. If live migration is required, select both Full Data Migration and Incremental Data Migration. <p> 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.</p>

Item	Description
Migration object	<ul style="list-style-type: none"> Select the database to be migrated from the Available list and click  to move it to the Selected list. <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p> Note: Data in the admin and local database is not migrated even if the admin database is selected.</p> </div> <ul style="list-style-type: none"> The migration object can be a database or a collection/function. By default, the object name will remain the same after it is migrated. If the object you migrate has a different name in the ApsaraDB for MongoDB instance, you must use the object name mapping feature provided by DTS. For more information, see Object name mapping.

8. After you configure the preceding settings, click **Precheck**.



Note:

- A pre-check must be performed before a migration task is formally started. The migration task can be started only after a successful pre-check.
- If the pre-check fails, click the  icon after each check item to view the failure details. Perform the pre-check again after you have rectified the failed items.

9. After the pre-check is successful, click **Next**.

10. On the **Confirm Settings** page, select **Channel Specification** and **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 full data migration task. Otherwise, the system may fail to migrate the full data. Wait until the migration task ends automatically.

- Incremental data migration

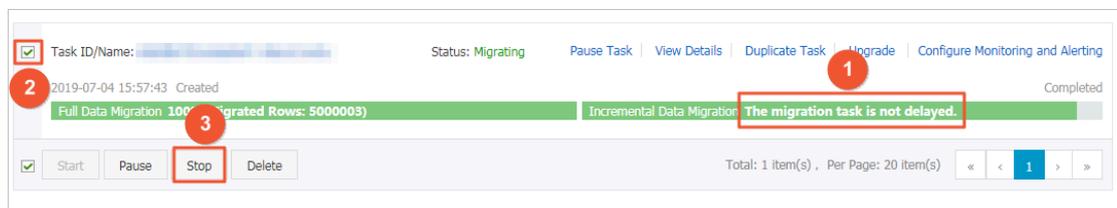
An incremental data migration task does not end automatically. You must manually stop the migration task.



Note:

Select an appropriate time to manually stop the migration task, such as during off-peak hours or when you are ready to switch your business to an ApsaraDB for MongoDB instance.

- a. When **Incremental Data Migration** and **The migration task is not delayed.** are displayed in the progress bar of the migration task, stop writing data to the source database for a few minutes. At this time, the latency time of **Incremental Data Migration** may be displayed in the progress bar of the migration task.
- b. When **Incremental Data Migration** and **The migration task is not delayed.** are displayed again in the progress bar of the migration task, manually stop the migration task.



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

8 Migrate data between instances under the same Alibaba Cloud account

8.1 Migrate data between RDS instances

This topic describes how to use Data Transmission Service (DTS) to migrate data between RDS instances. DTS supports schema migration, full data migration, and incremental data migration. When configuring a data migration task, you can select all of the supported migration types to ensure service continuity.

Prerequisites

The source and destination RDS instances must use the same database engines. DTS does not support data migration between RDS instances that use different database engines.

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%.
- 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.
- To ensure data consistency, do not write new data into the source RDS instance during full data migration.
- 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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

Migration types

- Schema migration

DTS migrates the schemas of the required objects to the destination instance.

- Full data migration

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

- Incremental data migration

After full data migration is complete, DTS synchronizes incremental data from the source RDS instance to the destination RDS instance. Incremental data migration helps you 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

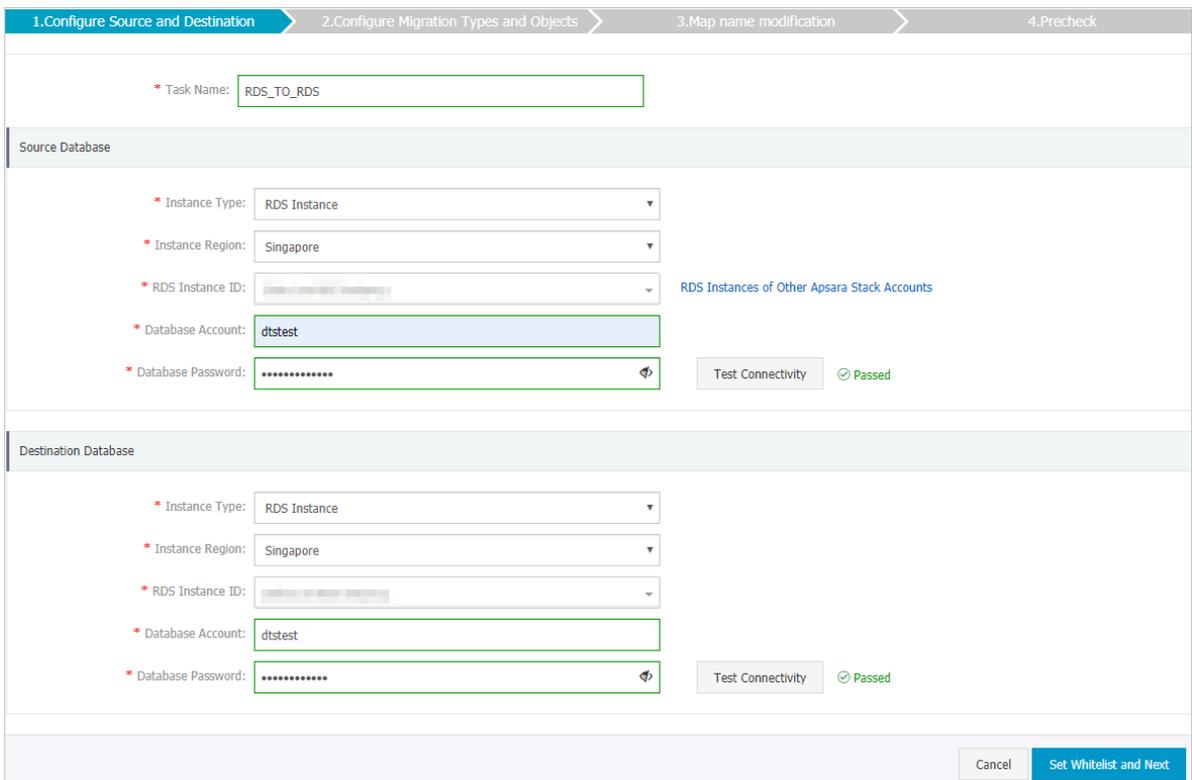
Procedure

1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

- At the top of the **Migration Tasks** page, select the region where the destination RDS instance resides.



- In the upper-right corner of the page, click **Create Migration Task**.
- 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.
Source Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination RDS instance resides.

Section	Parameter	Description
	RDS Instance ID	<p>Select the ID of the source RDS instance.</p> <p> 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.</p>
	Database Name	<p>Enter the name of the source database in the ApsaraDB RDS for PostgreSQL instance.</p> <p> Note: This parameter is required only if the database engine of the RDS instance is PostgreSQL.</p>
	Database Account	<p>Enter the database account of the source RDS instance. For more information about permissions required for the account, see Permissions required for database accounts.</p>
	Database Password	<p>Enter the password for the database account.</p> <p> Note: After the source database parameters are specified, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are correct. If the source database parameters are correct, the Test Passed message is displayed, If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database parameters as prompted.</p>
	Encryption	<p>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.</p> <p> Note: This parameter is required only if the database engine of the RDS instance is MySQL.</p> <p>The Encryption parameter is available only in mainland China and Hong Kong(China).</p>

Section	Parameter	Description
Destination Database	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.  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 Name	Enter the name of the destination database in the ApsaraDB RDS for PostgreSQL instance. The name of the destination database can be different from that of the source database.  Note: This parameter is required only if the database engine of the RDS instance is PostgreSQL .
	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 .
	Database Password	Enter the password for the database account.  Note: After the destination database parameters are specified, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are correct. If the destination database parameters are correct, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database parameters as prompted.

Section	Parameter	Description
	Encryption	<p>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.</p> <p> Note:</p> <p>This parameter is required only if the database engine of the RDS instance is MySQL.</p> <p>The Encryption parameter is available only in mainland China and Hong Kong(China).</p>

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 and destination RDS instances. This ensures that DTS servers can connect to the RDS instances.

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

The screenshot displays the configuration interface for data migration. At the top, there are four steps: 1. Configure Source and Destination, 2. Configure Migration Types and Objects (current step), 3. Map name modification, and 4. Precheck. Under 'Migration Types', three options are checked: Schema Migration, Full Data Migration, and Incremental Data Migration. A note explains that data migration applies to short-term scenarios like migrating data to the cloud or scaling/sharding databases. The main area is split into 'Available' and 'Selected' panels. The 'Available' panel shows a tree structure of database objects including 'mysqltest', 'Tables' (with 'customer2'), 'Views', 'Functions', 'Procedures', and 'sys'. The 'Selected' panel shows 'mysqltest (2Objects)' containing 'customer' and 'vipinfo'. Below the panels, there is a 'Name batch change' section with radio buttons for 'No' (selected) and 'Yes'. An 'Information' section provides two key points: 1. Data migration only copies data and schema from source to destination. 2. DDL operations are not supported during migration. At the bottom right, there are buttons for 'Cancel', 'Previous', 'Save', and 'Precheck'.

Parameter	Description
Migration types	<p>Select the migration types based on your requirements. The migration types must be supported by the database engine.</p> <ul style="list-style-type: none"> 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. <p>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.</p>

Parameter	Description
Objects to be migrated	<p>In the Available section, select the objects to be migrated and click the  icon to add the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> You can select databases, tables, or columns 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 source RDS instance. If the migrated object has a different name in the destination RDS 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. Click **Precheck** in the lower-right corner of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  corresponding to each failed item to view the details.

Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **Data Transmission Service (Pay-As-You-Go) Service Terms**.

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

- Full data migration

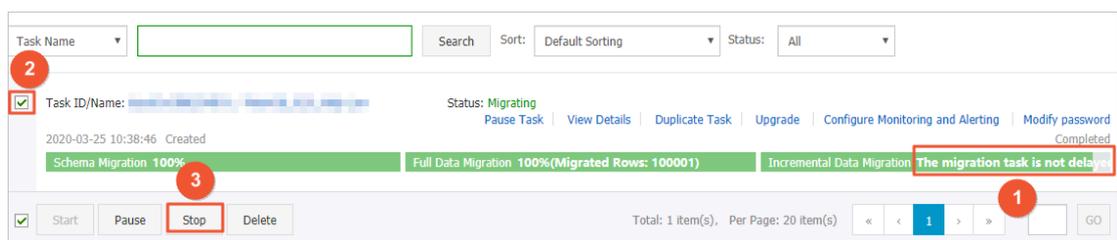
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination instance.

- a. When the task progress bar switches to **Incremental Data Migration** and **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



8.2 Migrate full data between ApsaraDB RDS for MySQL instances on a regular basis

This topic uses full data migration between ApsaraDB RDS for MySQL instances as an example to describe how to configure a scheduling policy for regular full data migration.

 **Note:**
All databases that support full data migration support the scheduling feature. For more information, see [Database and migration types](#).

Prerequisites

The source and destination ApsaraDB RDS for MySQL instances reside in the China (Qingdao), China (Beijing), China (Zhangjiakou-Beijing Winter Olympics), China (Shenzhen), or China (Shanghai) region. The scheduling feature for regular full data migration is available only in these regions. For more information, see [Create an ApsaraDB RDS for MySQL instance](#).

Background information

Scheduling is an advanced feature of data migration. You can configure a scheduling policy and scheduling cycle when you migrate schemas and historical data from the source database to the destination database on a regular basis. The scheduling feature allows you to flexibly build data warehouses. For more information, see [Scheduling policies of regular data migration](#).

Precautions

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.

Limits

- DTS supports schema migration of the following objects: tables, views, triggers, stored procedures, and stored functions.



Note:

During schema migration, the DEFINER mode of views, stored procedures, and stored functions is shifted to the INVOKER mode.

- The information of the source database account cannot be migrated. If you need to use views, stored procedures, and stored functions, you must grant read and write permissions to the destination database account.

Permissions required for database accounts

Database	Schema migration	Full data migration
Source RDS instance	The read/write permissions	The read/write permissions
Destination RDS instance	The read/write permissions	The read/write permissions



Note:

For more information about how to create and authorize a database account, see [Create an account for an ApsaraDB RDS for MySQL instance](#).

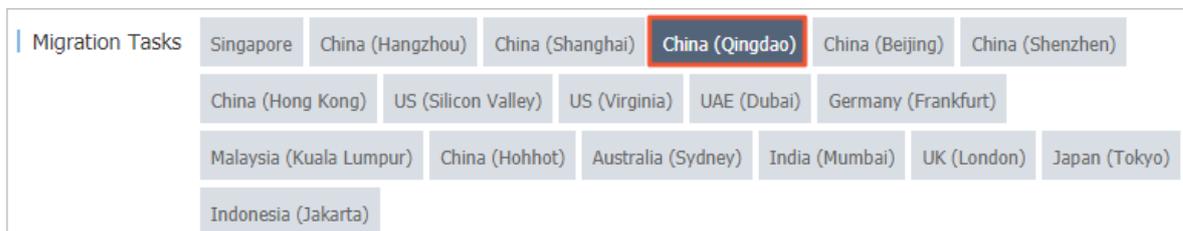
Billing

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 #unique_51 .

Procedure

The procedure in this topic uses ApsaraDB RDS for MySQL instances that reside in the China (Qingdao) region as an example.

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 **China (Qingdao)** region.



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.
Source Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the source RDS instance resides. In this example, select the China (Qingdao) region.
	Database 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 .

Section	Parameter	Description
	Database Password	<p>Enter the password of the source database account.</p> <p> Note: After you specify the source 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 source database parameters based on the check results.</p>
	Encryption	<p>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.</p> <p> Note: The Encryption parameter is available only for regions in mainland China and the Hong Kong (China) region.</p>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the destination RDS instance resides. In this example, select the China (Qingdao) region.
	Database 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 .
	Database Password	<p>Enter the password of the destination database account.</p> <p> 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.</p>

Section	Parameter	Description
	Encryption	<p>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.</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p> Note: The Encryption parameter is available only for regions in mainland China and the Hong Kong (China) region.</p> </div>

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

7. Configure a scheduling policy and select the objects to be migrated.

a) Configure a scheduling policy.

1.Configure Source and Destination
2.Configure Migration Types and Objects
3.Advanced Settings
4.Precheck

* Migration Types: Schema Migration Full Data Migration

During full data migration, data updates in the source database are not migrated to the destination instance. For data consistency, we recommend that you select Schema Migration, Full Data Migration, and Incremental Data Migration. When the CPU of the target instance MySQL is lower than 2 cores, it is recommended to close the slow log. [View More](#)

Free experience of periodic full-scale scheduling function: Beta ?

* Scheduler strategy: Duplicate mode, create new full replica to target [View the restrictions of each scheduling policy](#)

Large consumption of resources, you need to pay attention to the library table name changes

* Dispatch effective date: May 27, 2020 - May 30, 2020 ☰ ?

* Scheduling cycle: Days

* Specify a specific time: ☑ 27th ☑ 28th Can Multiple Choice

Task start time: 17 : 05

* Task failure policy: After a single task fails, subsequent scheduling tasks will not be exec [View Scheduled Scheduling Time](#)

- May 27, 2020, 17:05:00
- May 28, 2020, 17:05:00

Parameter	Description
Migration Types	<p>Select Schema Migration and Full Data Migration.</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p> Note: Do not select Incremental Data Migration. Otherwise, the Enable Scheduling for Regular Full Data Migration switch is not displayed.</p> </div>
Free experience of periodic full-scale scheduling function	Turn on this switch to enable the scheduling feature.

Parameter	Description
Scheduler strategy	<ul style="list-style-type: none">• Duplicate mode• Resynchronize mode• Based on the incremental mode of time stamp <p> Note: For more information about the scheduling policies, see Scheduling policies.</p>
Dispatch effective date	Specify the start date and end date of the schedule. The maximum effective duration is 365 days.
Scheduling cycle	Select Hour , Day , or Week as the scheduling cycle and set the specific scheduled time. <p> Note:</p> <ul style="list-style-type: none">• During full data migration, DTS uses read and write resources of the source and destination databases. This may increase the database load. Before you start scheduling, evaluate the performance of the source and destination databases. We recommend that you set the scheduled time to off-peak hours.• After you complete the settings, click Task Failure Policy next to View Scheduled Scheduling Time to check whether the set scheduled time meets the expectation.

Parameter	Description
Task Failure Policy	Select one of the following policies based on your business requirements: <ul style="list-style-type: none">• After a single task fails, subsequent scheduling tasks will not be executed• After a single task fails, ignore the failure and execute the next scheduled task

b) Select the objects to be migrated.

Available

Expand the tree before you perform a glob | Q

- sys
- dtstestdata
- mysql_xtdkamka
- dts

Select All

Selected (To edit an object name or its filter, hover over the object and click Edit.) [Learn more.](#)

dtstestdata (10Objects)

- order

Remove All

Change Mapped Name: Do Not Change Database and Table Names Change Database and Table Names

Information:

1. Data migration only copies the data and schema in the source database and saves the copy in the destination database. The process does not affect any data or schema in the source database.
2. DDL operations are not supported during data migration because this can cause migration failures.

Cancel Previous Advanced Settings Save Precheck

Parameter	Description
Objects to be migrated	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> You can select columns, tables, or databases as the objects to be migrated. After an object is migrated to the destination instance, the name of the object remains unchanged. You can change the names of the objects that are migrated to the destination 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. </div>

c) In the lower-right corner of the page, click **Advanced Settings**. On the page that appears, specify the timestamp field.

 **Note:**

This step is required only if you set **Scheduling Policy** to **Timestamp-based Full Data Mode**.

1.Configure Source and Destination > 2.Configure Migration Types and > **3.Advanced Settings** > 4.Precheck

Database name	Table Name	Timestamp field
dtstestdata	order	ordertime

Total: 1 item(s), Per Page: 20 item(s) << < 1 > >>

Cancel Previous Save **Precheck**

8. Click **Precheck** on the lower right of the page.

 **Note:**

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.
10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **Data Transmission Service (Pay-As-You-Go) Service Terms**.
11. Click **Buy and Start**. The migration task is started based on the specified scheduled time.

<input type="checkbox"/>	Task ID/Name: [redacted]	Status: Passed	Stop Regular Full Data Migration
	2020-05-19 14:45:31 Created	Full scheduler: RUNNING	Dispatch effective date: 2020-05-20~2020-05-26 Next Scheduler Time: 2020-05-20 15:45:00
	Schema Migration 0%	Full Data Migration 0%	
<input type="checkbox"/>	Start	Pause	Stop Delete
Total: 1 item(s), Per Page: 20 item(s)			<< < 1 > >>

**Note:**

The migration task automatically stops at the end of the scheduled time. To stop the task before the end time of the schedule, click **Stop Regular Full Data Migration**.

Result

If a full data migration task is completed before the end of the current scheduled time, the task status changes to **Completed**. DTS waits for the next scheduled time and performs full data migration again.

Task ID/Name: [redacted]	Status: Migrating	Stop Regular Full Data Migration
May 27, 2020, 17:11:00 Created	Full scheduler: RUNNING	Dispatch effective date: May 27, 2020~May 30, 2020 Next Scheduler Time: May 28, 2020, 17:11:00
Schema Migration 0%	Full Data Migration 0%	

8.3 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 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 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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.
- 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 data migration

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



Note:

- 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.
- 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

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul style="list-style-type: none"> • 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 **Migration Tasks** the 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.

The screenshot shows the '1. Configure Source and Destination' step of a DTS migration task. The task name is 'MariaDB_TO_MYSQL'. The source database is configured with Instance Type: RDS Instance, Instance Region: China (Hangzhou), RDS Instance ID: rm-xxxxxxqc, Database Account: dtstest, and Database Password: *****. The destination database is configured with Instance Type: RDS Instance, Instance Region: China (Hangzhou), RDS Instance ID: rm-xxxxxxn7, Database Account: dtstest, and Database Password: *****. Both source and destination connectivity tests have passed. The encryption option is set to 'Non-encrypted'. Buttons for 'Cancel' and 'Set Whitelist and Next' are visible at the bottom right.

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 Database	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.
	Database 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 .

Section	Parameter	Description
	Database Password	<p>Enter the password of the database account.</p> <p>After you specify the source database parameters, click Test Connectivity next to the Database Password parameter to verify whether the parameters are valid.</p> <div style="background-color: #f0f0f0; padding: 5px;">  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. </div>
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.
	Database 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	<p>Enter the password of the database account.</p> <p>After you specify the destination database parameters, click Test Connectivity next to the Database Password parameter to verify whether the parameters are valid.</p> <div style="background-color: #f0f0f0; padding: 5px;">  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. </div>

Section	Parameter	Description
	Encryption	<p>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.</p> <p> Note: The Encryption parameter is available only for data migration instances that reside in mainland China and Hong Kong (China).</p>

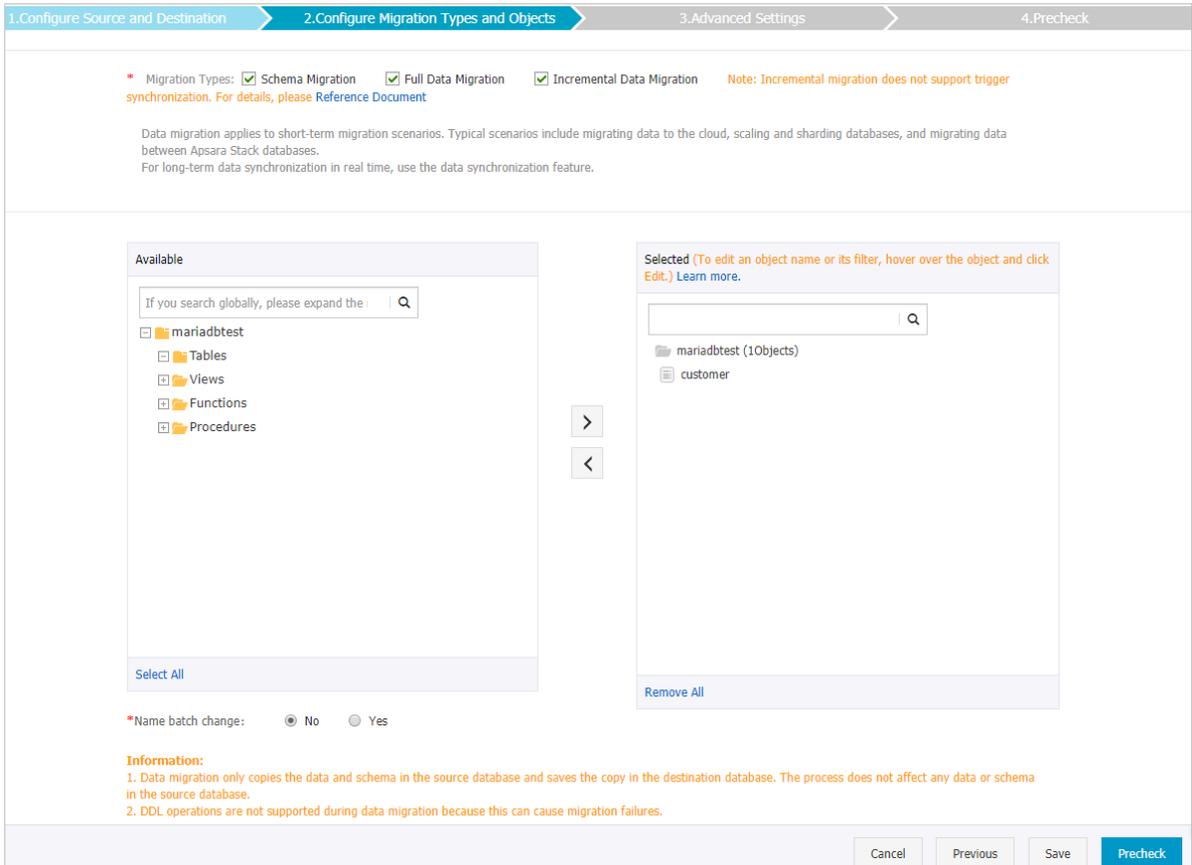
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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.



Item	Description
Migration types	<ul style="list-style-type: none"> To perform only full data migration, select Schema Migration and Full Data Migration. If you want to migrate data without business disruptions, select Schema Migration, Full Data Migration, and Incremental Data Migration. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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 [#unique_59](#). If you do not need to switch your workloads, you can stop the migration task by using the following procedure.

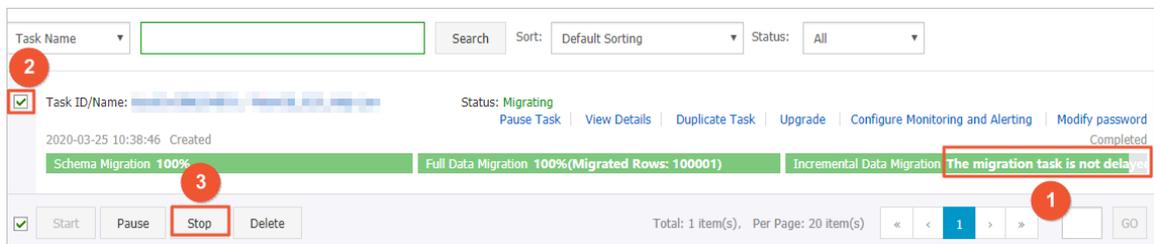
- Full data migration

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

- Incremental data migration

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

1. 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**.
2. 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.

8.4 Migrate data from an ApsaraDB RDS for MySQL database to an ApsaraDB for PolarDB database

ApsaraDB for 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. You can use Data Transmission Service (DTS) to migrate data from an ApsaraDB RDS for MySQL database to an ApsaraDB for PolarDB database.

Prerequisites

An ApsaraDB for PolarDB cluster is created. For more information, see [Create an ApsaraDB for PolarDB cluster](#).

Background information

- 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%.
- If your ApsaraDB RDS for MySQL database does not have primary key or unique constraints and each field in the database has duplicate values, the data migrated to the destination database may be duplicated.
- Concurrent insertions are performed during full data migration. This results in table fragmentation in the destination instance. After a full data migration task is completed, the tablespace of the destination instance is larger than that of the source instance.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the 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 of the following objects: tables, views, triggers, stored procedures, and stored functions.

**Note:**

During schema migration, the DEFINER mode of views, stored procedures, and stored functions is shifted to the INVOKER mode.

- The information of the source database account cannot be migrated. If you need to use views, stored procedures, and stored functions, you must grant read and write permissions to the destination database account.

Migration types

DTS supports schema migration, full data migration, and incremental data migration. For more information, see [#unique_73](#).



Note:

You can use these three migration types together to migrate data without service interruptions.

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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

SQL operations that can be synchronized during incremental data migration

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul style="list-style-type: none"> 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 permissions
ApsaraDB RDS for MySQL	Read permission on objects to be migrated

Database	Required permissions
ApsaraDB for PolarDB	Read and write permissions on migrated objects

**Note:**

For more information about how to create and authorize a database account, see [Create an ApsaraDB RDS for MySQL database account](#) and [Create an ApsaraDB for PolarDB 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 **Migration Tasks** the 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.

The screenshot shows the '1. Configure Source and Destination' step of a DTS migration task. The task name is 'RDS_TO_PolarDB'. The source database is an RDS Instance in the Singapore region with ID [redacted] and account 'dtstest'. The destination database is a PolarDB instance in the Singapore region with ID [redacted] and account 'dtstest'. Both source and destination connectivity tests have passed.

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.
Source Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the RDS instance resides.
	Database Account	Enter the ApsaraDB RDS for MySQL database account. For more information about the permissions required for the ApsaraDB RDS for MySQL database account, see Permissions required for database accounts .

Section	Parameter	Description
	Database Password	<p>Enter the password of the ApsaraDB RDS for MySQL database account.</p> <p> Note: After you specify the source database parameters, click Test Connectivity next to the Database Password parameter to verify whether the parameters are valid. If the source database parameters are valid, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database parameters as prompted.</p>
	Encryption	<p>Select Non-encrypted or SSL-encrypted as needed. If you select SSL-encrypted, you must enable the SSL encryption feature for the RDS instance. For more information about how to enable the feature, see Configure SSL encryption for a RDS instance.</p> <p> Note: Encryption is available only in mainland China and Hong Kong(China).</p>
Destination Database	Instance Type	Select PolarDB .
	Instance Region	Select the region where the ApsaraDB for PolarDB cluster resides.
	PolarDB Instance ID	Enter the ApsaraDB for PolarDB cluster ID.
	Database Account	Enter the ApsaraDB for PolarDB database account. For more information about the permissions required for the ApsaraDB for PolarDB database account, see Permissions required for database accounts .

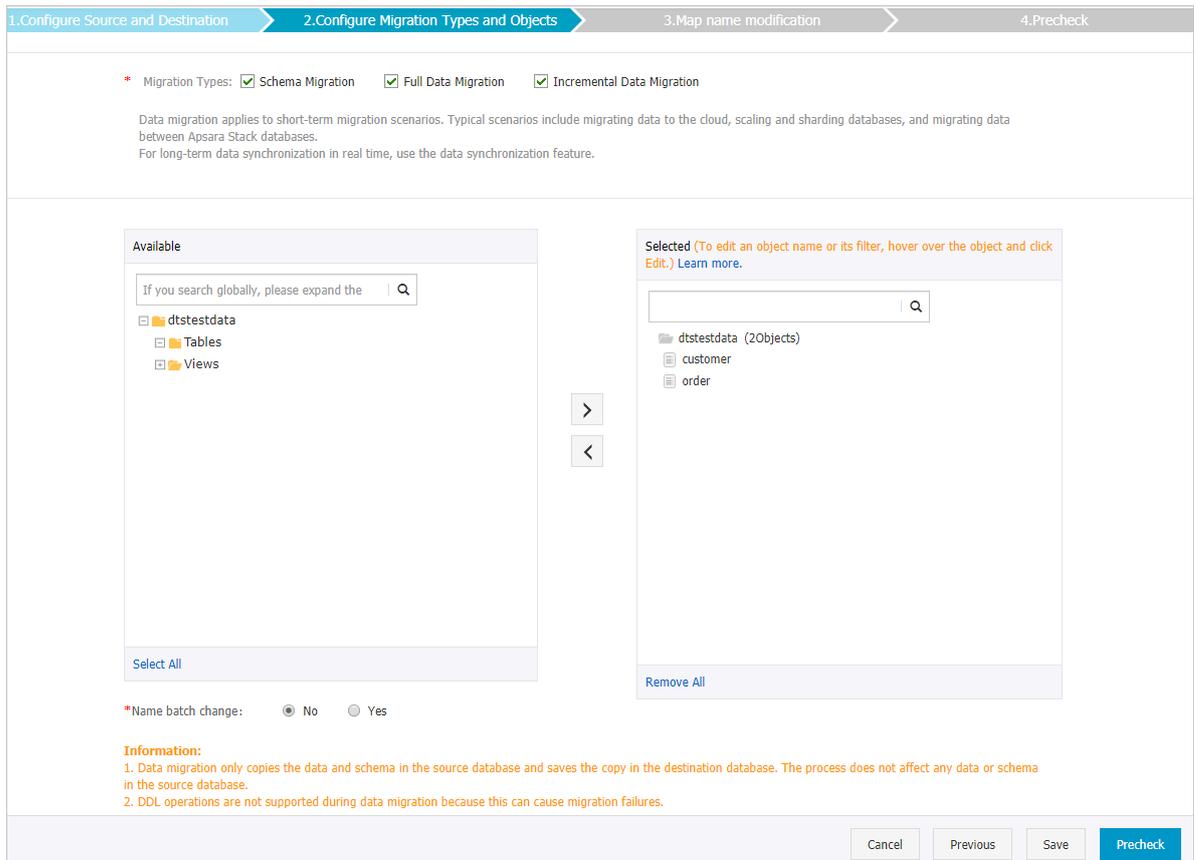
Section	Parameter	Description
	Database Password	<p>Enter the password of the ApsaraDB for PolarDB database account.</p> <p> Note: After you specify the destination database parameters, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are valid. If the destination database parameters are valid, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database parameters as prompted.</p>

6. Click **Set Whitelist and Next** in the lower-right corner of the page.

**Note:**

The IP addresses of DTS servers are added to the whitelist of the ApsaraDB for PolarDB cluster and RDS instance. Then, DTS servers can connect to the cluster and RDS instance.

7. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> To perform only full data migration, select Schema Migration and Full Data Migration. If you want to migrate data without business disruptions, select Schema Migration, Full Data Migration, and Incremental Data Migration. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

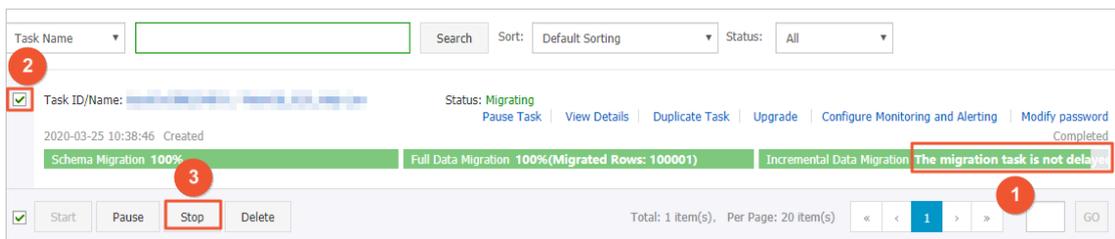
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When 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 cluster.

8.5 Migrate data between ApsaraDB for PolarDB clusters

ApsaraDB for 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. You can use Data Transmission Service (DTS) to migrate data between ApsaraDB for PolarDB clusters.

Prerequisites

- An ApsaraDB for PolarDB cluster is created. For more information, see [Create an ApsaraDB for PolarDB cluster](#).
- The binary logging feature for the ApsaraDB for PolarDB cluster is enabled. For more information, see [Enable binlogging](#).

Background information

- 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%.
- If your ApsaraDB RDS for MySQL database does not have primary key or unique constraints and each field in the database has duplicate values, the data migrated to the destination database may be duplicated.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the 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 [#unique_73](#).

**Note:**

You can use these three migration types together to migrate data without service interruptions.

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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

SQL operations that can be synchronized during incremental data migration

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul style="list-style-type: none"> 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 permissions	
Source ApsaraDB for PolarDB cluster	Read permission on objects to be migrated	
Destination ApsaraDB for PolarDB cluster	Read and write permissions on migrated objects	

**Note:**

For more information about how to create and authorize a database account, see [Create an ApsaraDB RDS for MySQL database account](#).

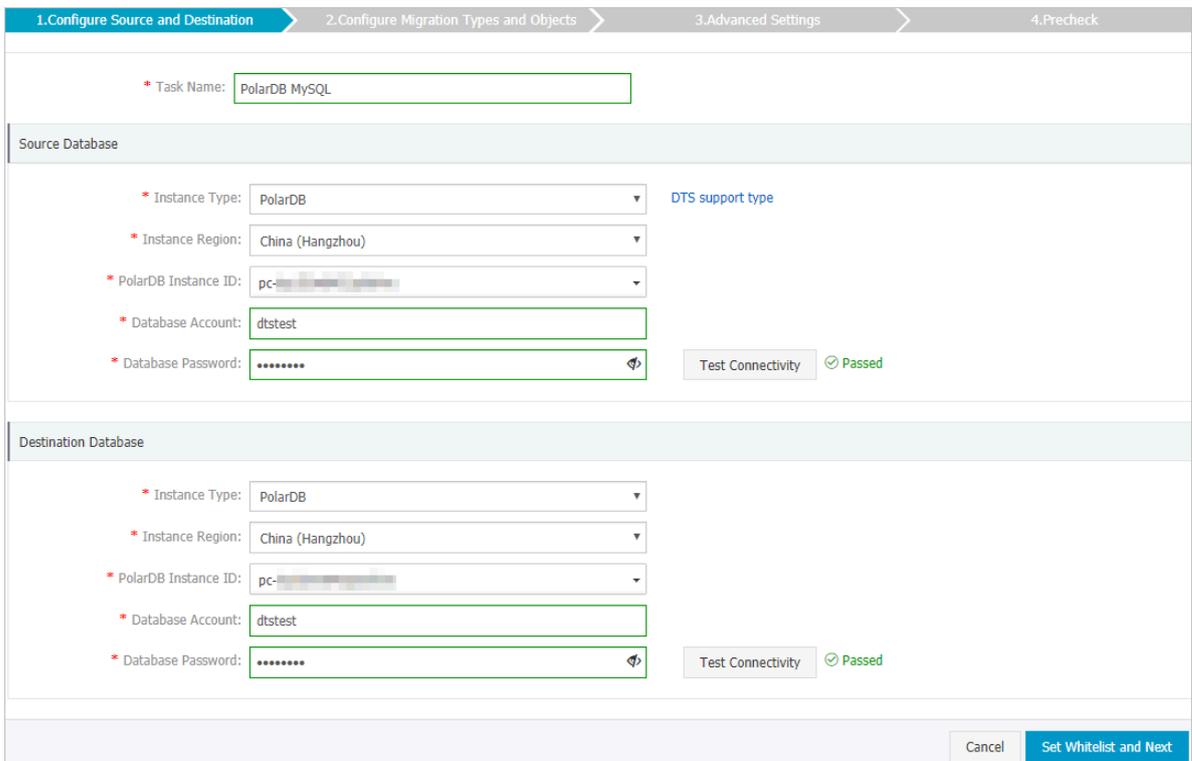
Procedure

1. Log on to the [DTS console](#).

- In the left-side navigation pane, click **Data Migration**.
- At the top of **Migration Tasks** the page, select the region where the destination cluster resides.



- In the upper-right corner of the page, click **Create Migration Task**.
- 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.
Source Database	Instance Type	Select PolarDB .
	Instance Region	Select the region where the source ApsaraDB for PolarDB cluster resides.
	PolarDB Instance ID	Select the source ApsaraDB for PolarDB cluster ID.
	Database Account	Enter the database account of the source ApsaraDB for PolarDB cluster.

Section	Parameter	Description
	Database Password	<p>Enter the password of the database account.</p> <p> Note: After the source database parameters are specified, click Test Connectivity next to the Database Password parameters to verify whether the specified parameters are correct. If the source database parameters are correct, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database parameters as prompted.</p>
Destination Database	Instance Type	Select PolarDB .
	Instance Region	Select the region where the destination ApsaraDB for PolarDB cluster resides.
	PolarDB Instance ID	Enter the destination ApsaraDB for PolarDB cluster ID.
	Database Account	Enter the database account of the destination ApsaraDB for PolarDB cluster.
	Database Password	<p>Enter the password of the database account.</p> <p> Note: After the destination database parameters are specified, click Test Connectivity next to the Database Password parameter to verify whether the specified parameters are correct. If the destination database parameters are correct, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database parameters as prompted.</p>

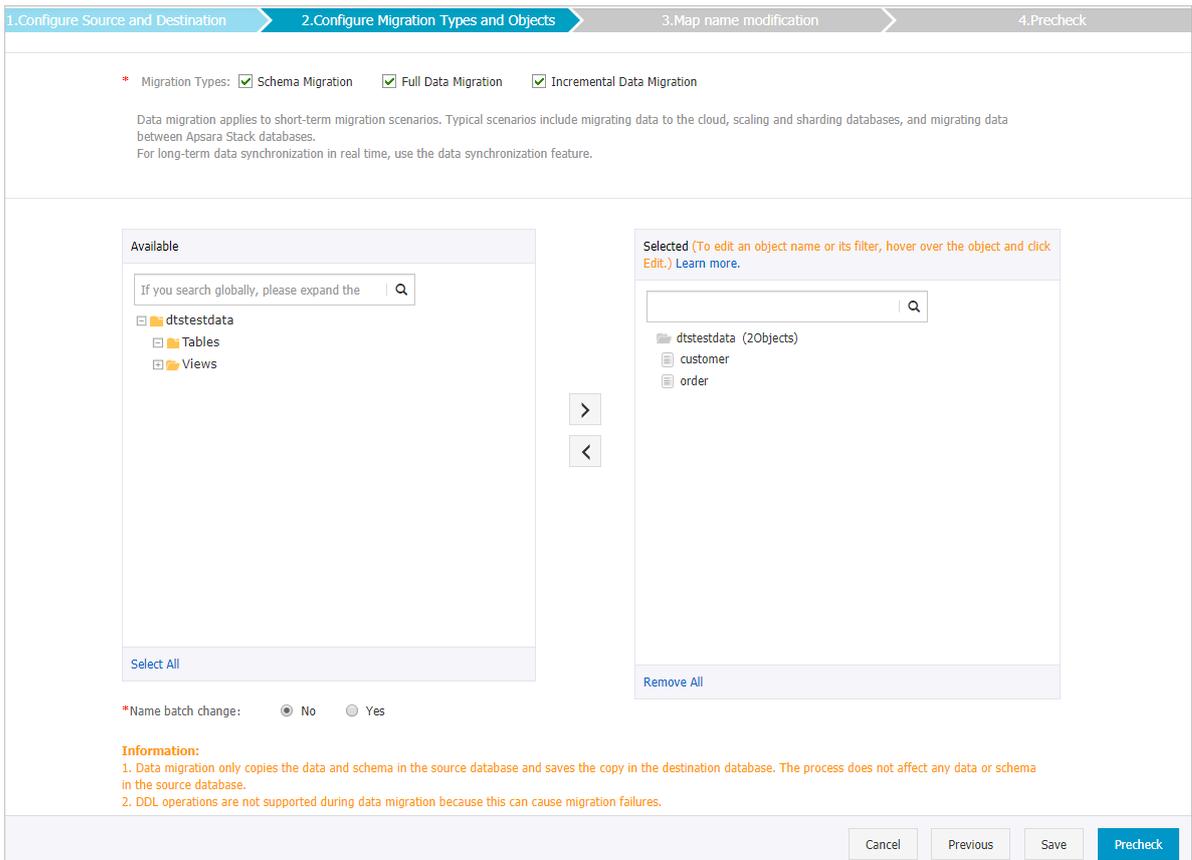
- Click **Set Whitelist and Next** in the lower-right corner of the page.



Note:

The IP addresses of DTS servers are added to the whitelist of the source and destination ApsaraDB for PolarDB clusters. This makes sure that the DTS server can connect to the clusters.

7. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> To perform only full data migration, select Schema Migration and Full Data Migration. If you want to migrate data without business disruptions, select Schema Migration, Full Data Migration, and Incremental Data Migration. <p> 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.</p>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

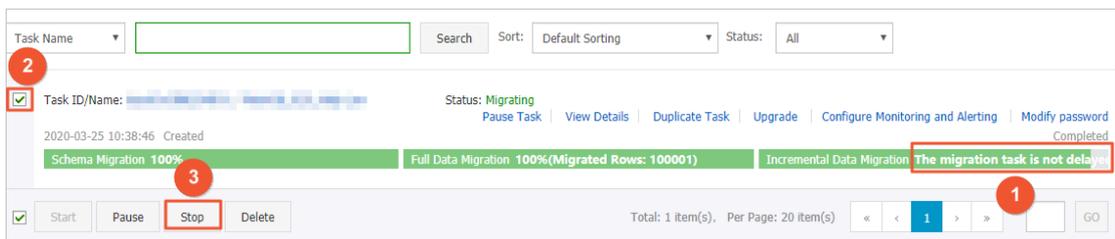
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When 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 cluster.

What's next

After the data migration is complete, you must delete the accounts of both the source and destination databases to ensure security.

8.6 Migrate data from an ApsaraDB for PolarDB database to an ApsaraDB RDS for MySQL database

ApsaraDB RDS for MySQL is a reliable and elastic online database service provided by Alibaba Cloud. It is a complete solution that can be used to implement disaster recovery,

data backup, data recovery, and data migration. You can use Data Transmission Service (DTS) to migrate data from an ApsaraDB for PolarDB database to an ApsaraDB RDS for MySQL database.

Prerequisites

- The binary log feature for the ApsaraDB for PolarDB cluster is enabled. 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](#).

Background information

- 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%.
- If your ApsaraDB RDS for MySQL database does not have primary key or unique constraints and each field in the database has duplicate values, the data migrated to the destination database may be duplicated.
- Concurrent insertions are performed during full data migration. This results in table fragmentation in the destination instance. After a full data migration task is completed, the tablespace of the destination instance is larger than that of the source instance.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the 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 of the following objects: tables, views, triggers, stored procedures, and stored functions.



Note:

During schema migration, the DEFINER mode of views, stored procedures, and stored functions is shifted to the INVOKER mode.

- The information of the source database account cannot be migrated. If you need to use views, stored procedures, and stored functions, you must grant read and write permissions to the destination database account.

Migration types

DTS supports schema migration, full data migration, and incremental data migration. For more information, see [#unique_73](#).



Note:

You can use these three migration types together to migrate data without service interruptions.

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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

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	<ul style="list-style-type: none"> 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 permissions
ApsaraDB for PolarDB	Read permission on objects to be migrated
ApsaraDB RDS for MySQL	Read and write permissions on migrated objects



Note:

For more information about how to create and authorize a database account, see [Create an ApsaraDB for PolarDB database account](#) and [Create an ApsaraDB RDS for MySQL 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 source and destination databases.

The screenshot shows the '1. Configure Source and Destination' step of a DTS migration task. The task name is 'PolarDB_To_RDS'. The source database is configured as a PolarDB instance in the China (Hangzhou) region with ID 'pc-...' and database account 'dtstest'. The destination database is configured as an RDS Instance in the China (Hangzhou) region with ID 'rm-...' and database account 'dtstest'. Both source and destination connectivity tests have passed. The encryption option is set to 'Non-encrypted'. Buttons for 'Cancel' and 'Set Whitelist and Next' are visible at the bottom right.

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.
Source Database	Instance Type	Select PolarDB .
	Instance Region	Select the region where the source ApsaraDB for PolarDB cluster resides.
	PolarDB Instance ID	Select the ApsaraDB for PolarDB cluster ID.
	Database Account	Enter the ApsaraDB for PolarDB database account. For more information about the permissions required for the ApsaraDB for PolarDB database account, see Permissions required for database accounts .

Section	Parameter	Description
	Database Password	<p>Enter the password of the ApsaraDB for PolarDB database account.</p> <div style="background-color: #f0f0f0; padding: 10px;">  Note: After you specify the source database parameters, click Test Connectivity next to the Database Password parameter to verify whether the parameters are valid. If the source database parameters are valid, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the source database parameters as prompted. </div>
Destination Database	Instance Type	Select RDS Instance .
	Instance Region	Select the region where the RDS instance resides.
	Database Account	Enter the ApsaraDB RDS for MySQL database account. For more information about the permissions required for the ApsaraDB RDS for MySQL database account, see Permissions required for database accounts .
	Database Password	<p>Enter the password of the ApsaraDB RDS for MySQL database account.</p> <div style="background-color: #f0f0f0; padding: 10px;">  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 destination database parameters are valid, the Test Passed message is displayed. If the Test Failed message is displayed, click Diagnose in the Test Failed message. Modify the destination database parameters as prompted. </div>

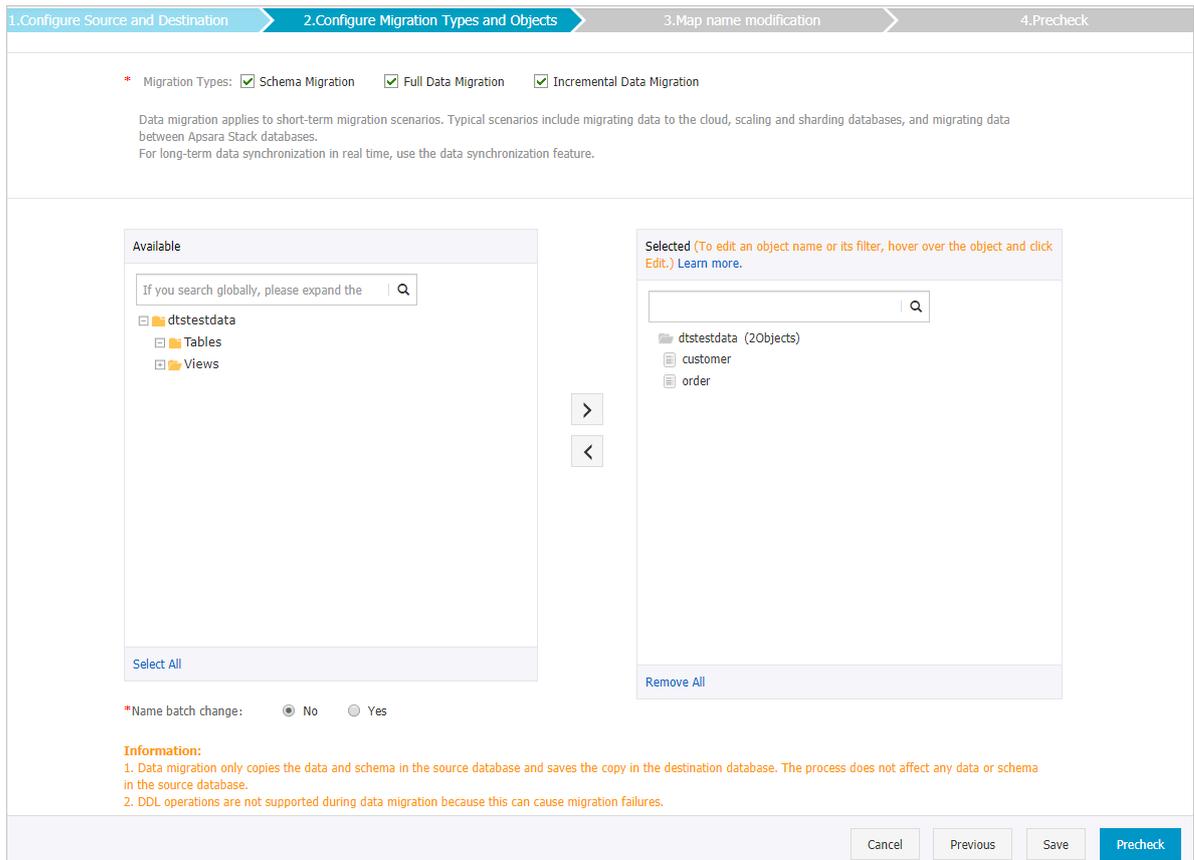
Section	Parameter	Description
	Encryption	<p>Select Non-encrypted or SSL-encrypted as needed. If you select SSL-encrypted, you must enable the SSL encryption feature for the RDS instance. For more information about how to enable the feature, see Configure SSL encryption for a RDS instance.</p> <p> Note: Encryption is available only in mainland China and Hong Kong(China).</p>

6. Click **Set Whitelist and Next** in the lower-right corner of the page.

**Note:**

The IP addresses of DTS servers are added to the whitelist of the ApsaraDB for PolarDB cluster and RDS instance. This makes sure that DTS servers can connect to the PolarDB cluster and RDS instance.

7. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

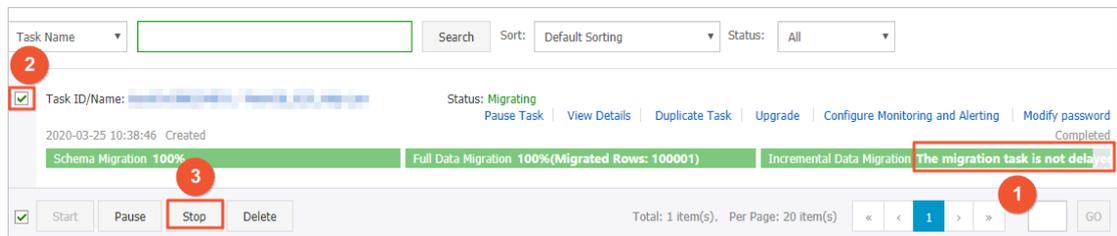
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When 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 RDS instance.

8.7 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 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%.

- 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 heartbeat table named `dts_postgres_heartbeat` to the source database. The following figure shows the schema of the heartbeat table.

SLOT_NAME	REVICE_TIME	REVICE_LSN	FLUSHED_LSN	UPDATE_TIME	DTS_SERVICE_TIME
w8i	1585104942560	0/44	null	2020-03-25 10:55:47.585187+08	1585104947579

- 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	Instance configurat ions	Internet traffic
Schema migration and full data migration	Free of charge.	Charged only when data is migrated from Alibaba Cloud over the Internet.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

Migration type	Description
Schema migration	<p>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.</p> <div style="border: 1px solid gray; padding: 5px; background-color: #f0f0f0;"> <p> Notice: However, if an object contains triggers, data will become inconsistent between the source and destination databases.</p> </div>

Migration type	Description
Full data migration	<p>DTS migrates historical data of the required objects from the source PolarDB cluster to the destination PolarDB cluster.</p> <p> 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.</p>
Incremental data migration	<p>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:</p> <ul style="list-style-type: none"> • INSERT, UPDATE, and DELETE operations • CREATE TABLE operations <p> Notice: The CREATE TABLE operations used to create partition tables or tables that contain functions cannot be synchronized.</p> <ul style="list-style-type: none"> • ALTER TABLE, DROP TABLE, RENAME TABLE, CREATE INDEX, and ADD INDEX operations <p>Incremental data migration allows you to ensure service continuity when you migrate data between PolarDB clusters.</p>

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 [#unique_93](#).



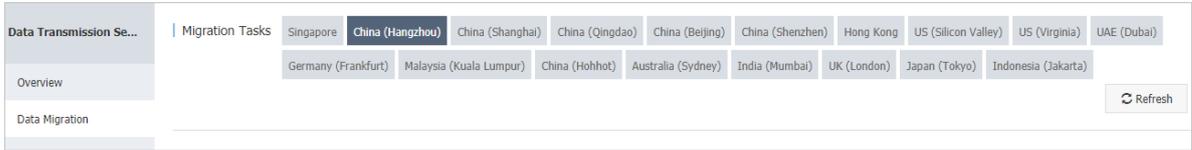
Warning:

If you change the value of the wal_level parameter, the source PolarDB cluster is restarted. Perform this operation with caution.

Procedure

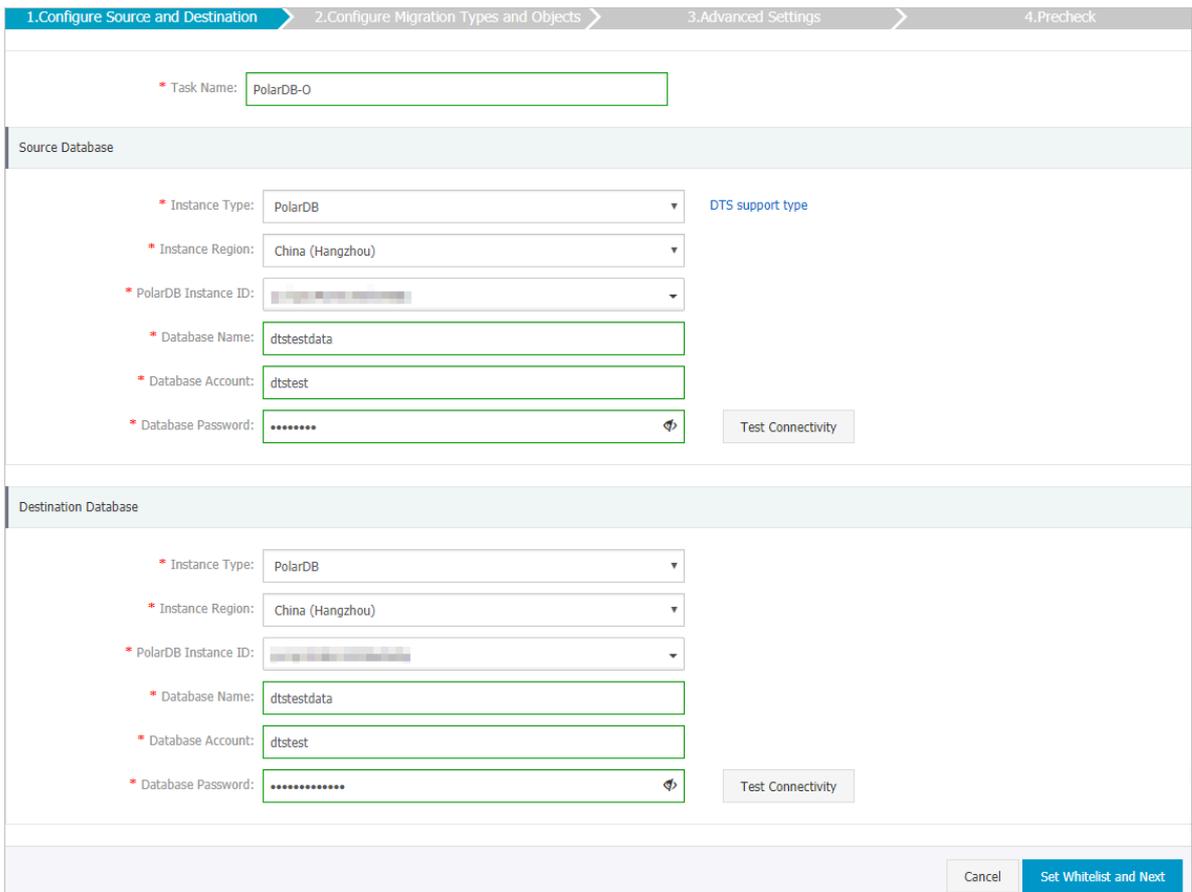
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of **Migration Tasks** the 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.



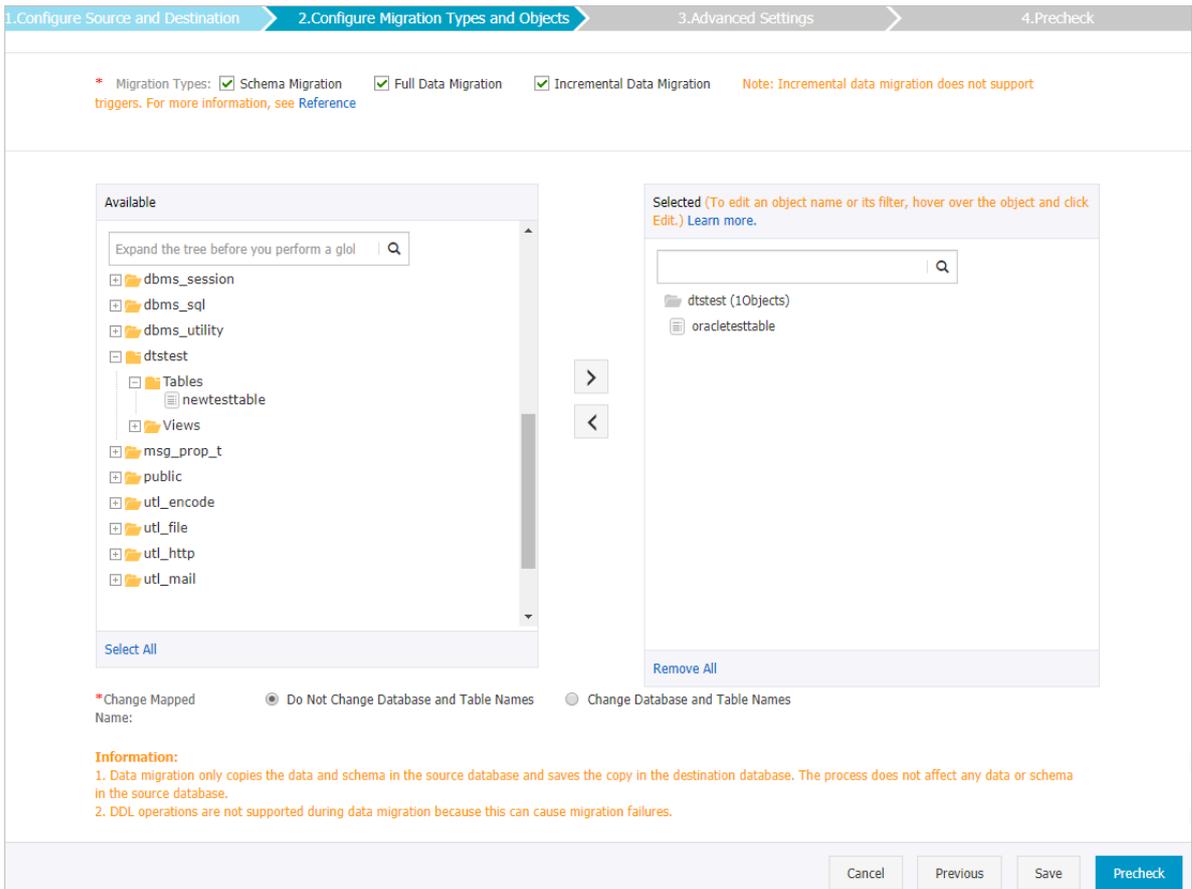
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	Select PolarDB .
	Instance Region	Select the region where the source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

Section	Parameter	Description
	Database Name	Enter the name of the source database.
	Database Account	Enter the privileged account of the source PolarDB cluster. For more information, see #unique_94 .
	Database 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.
Destination 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.
	Database Name	Enter the name of the destination database.
	Database Account	Enter the database account of the destination PolarDB cluster. The account must have the owner permission on the database .  Notice: You can specify the database owner when you create a 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 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**.

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
Migration Types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Notice:</p> <ul style="list-style-type: none"> 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. </div>

Parameter	Description
Objects	<p>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.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Notice:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.

 **Note:**

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

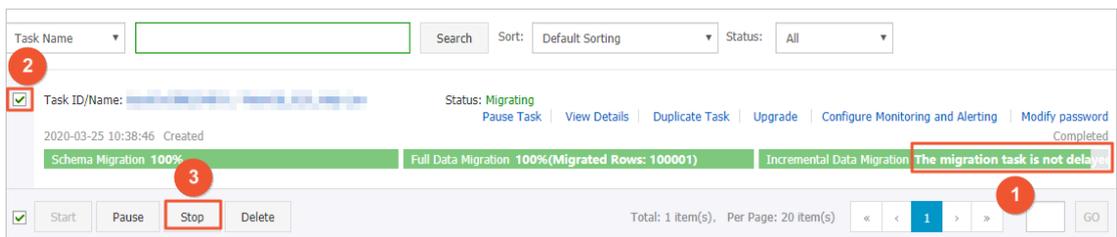
Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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

 **Note:**
Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



Related topics

[#unique_59](#)

8.8 Migrate data from a standalone instance to a replica set instance or a sharded cluster instance

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). A standalone instance is designed for development, testing, and other scenarios where non-

core enterprise data is stored. A replica set instance and a sharded cluster instance are more suitable for production scenarios .

Prerequisites

The storage capacity of the destination instance is greater than that of the source instance.

Precautions

- We recommend that you migrate your data during off-peak hours to avoid business interruptions.
- You can use DTS only to fully migrate data of a standalone instance. For data consistency, 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 [#unique_79](#).

Billing

Migration type	Link configuration fee	Internet traffic fee
Full data migration	Free of charge	Free of charge

Migration types

Full data migration is described in this topic. All data of the migration objects is migrated from the source instance to the destination instance.



Note:

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

Required database account permissions

Data source	Full data migration
Source ApsaraDB for MongoDB instance	Read permissions on the source database
Destination ApsaraDB for MongoDB instance	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](#).

Procedure

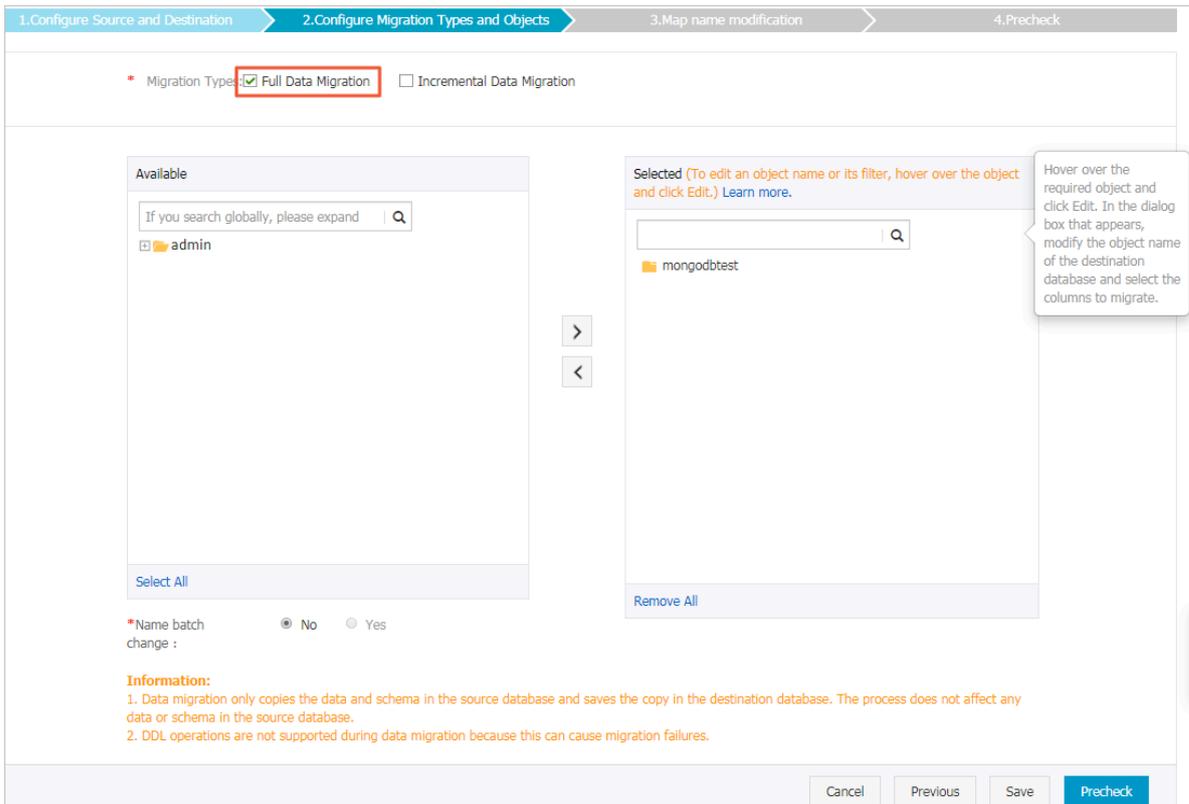
1. Configure both the source and destination databases.

The screenshot shows the '1. Configure Source and Destination' step of a DTS migration task. It features a progress bar at the top with four steps: '1. Configure Source and Destination' (active), '2. Configure Migration Types and Objects', '3. Map name modification', and '4. Precheck'. Below the progress bar, there is a 'Task Name' input field. The main configuration area is divided into two sections: 'Source Database' and 'Destination Database'. Each section contains the following fields: 'Instance Type' (dropdown), 'Instance Region' (dropdown), 'MongoDB Instance ID' (dropdown), 'Database Name' (text input), 'Database Account' (text input), and 'Database Password' (password input). A 'Test Connectivity' button is located to the right of the password field in each section. At the bottom right of the form, there are three buttons: 'Cancel', 'Assess Data Migration to Cloud', and 'Set Whitelist and Next'.

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify your own task name that helps identify the task. Task names do not need to be unique.
Source Database	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.
	Database Name	Enter the name of the authentication database. It is the database where the database account is created. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;">  Note: If the database account is root, enter admin. </div>

Section	Parameter	Description
	Database 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 .
	Database Password	Enter the password of the 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 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.
Destination Database	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.
	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 .
	Database 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 Test Failed message to modify the information as prompted.

2. Configure migration types and migration objects.



Parameter	Description
Migration Types	<p>Select Full Data Migration.</p> <p> Note: If the data source is a standalone instance, you can only select Full Data Migration. To ensure data consistency, do not write data to the source database while full data migration of a standalone instance is in progress.</p>
Available	<ul style="list-style-type: none"> In the Available section, select the objects you want to migrate and then click the  icon to move them to the Selected section. <p> Note: Data in the admin database cannot be migrated even if this database is selected.</p> <ul style="list-style-type: none"> 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.

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

Task ID/Name: XXXXXXXXXX Status: Completed
[View Details](#) | [Duplicate Task](#) | [Configure Monitoring and Alerting](#) | [Modify password](#)
Dec 16, 2019, 14:11:34 Created Dec 16, 2019, 14:13:49 Completed

Full Data Migration 100%(Migrated Rows: 677311)



Note:

Do not manually end a migration task. If you do so, the system may fail to migrate all data of the database. Wait until the migration task is complete.

4. Switch over your business to the destination ApsaraDB for MongoDB instance.

References

- [#unique_83](#)
- [#unique_95](#)
- [Configure sharding to maximize the performance of shards](#)

8.9 Migrate data from a replica set instance to a sharded cluster instance

This topic describes how to migrate data from a replica set instance to a sharded cluster instance 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 to ApsaraDB for MongoDB without interruptions to your business.

Prerequisites

Make sure that the shards in the destination sharded cluster instance have sufficient storage space.

Precautions

- We recommend that you migrate your data during off-peak hours to avoid business interruptions.
- 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 [#unique_79](#).

Billing

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

Migration types

Migration type	Description
Full data migration	<p>All data of the migration objects is migrated from the source instance to the destination instance.</p> <p> Note: Data migration is supported at the database, collection, and index levels.</p>
Incremental data migration	<p>Updated data of the migration objects is synchronized from the source instance to the destination instance.</p> <p> Note:</p> <ul style="list-style-type: none"> The create and delete operations on databases, collections, and indexes can be synchronized. The create, delete, and update operations on documents can be synchronized.

Required database account permissions

Data source	Full data migration	Incremental data migration
Source ApsaraDB for MongoDB replica set instance	Read permissions on the source database	Read permissions on the source database, admin database, and local database
Destination ApsaraDB for MongoDB sharded cluster 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

Create databases and collections to be sharded in the destination ApsaraDB for MongoDB instance, and configure data sharding based on your business needs. For more information, see [Configure sharding to maximize the performance of shards](#).



Note:

After data sharding is configured, data is no longer migrated to the same shard. This helps maximize cluster performance.

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 where the destination ApsaraDB for MongoDB instance resides.



4. In the upper-right corner, click **Create Migration Task**.

5. Configure both the source and destination databases.

The screenshot shows the configuration interface for a DTS task. It is divided into four steps: 1. Configure Source and Destination, 2. Configure Migration Types and Objects, 3. Map name modification, and 4. Precheck. The 'Task Name' field is at the top. Below are two sections: 'Source Database' and 'Destination Database'. Each section contains dropdown menus for Instance Type, Instance Region, and MongoDB Instance ID, and text input fields for Database Name, Database Account, and Database Password. A 'Test Connectivity' button is present next to the password fields. At the bottom right, there are buttons for 'Cancel', 'Assess Data Migration to Cloud', and 'Set Whitelist and Next'.

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify your own task name that helps identify the task. Task names do not need to be unique.
Source Database	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.
	Database Name	Enter the name of the authentication database. It is the database where the database account is created. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;">  Note: If the database account is root, enter admin. </div>
	Database 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 .

Section	Parameter	Description
	Database Password	<p>Enter the password of the database account.</p> <p> 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.</p>
Destination Database	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.
	Database Name	<p>Enter the name of the authentication database. It is the database where the database account is created.</p> <p> Note: If the database account is root, enter admin.</p>
	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 .
	Database Password	<p>Enter the password of the database account.</p> <p> 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.</p>

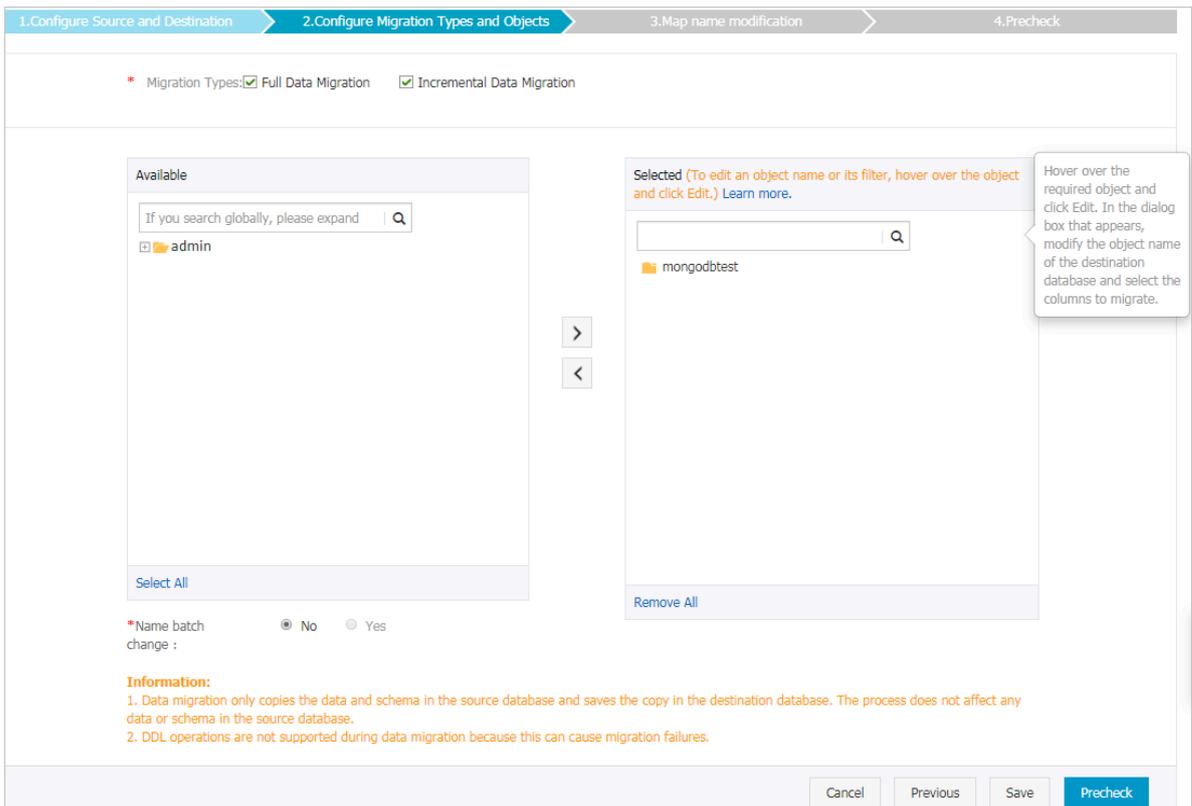
6. In the lower-right corner, click **Set Whitelist and Next**.



Note:

The IP addresses of DTS servers are automatically added to the whitelists of the source and destination ApsaraDB for MongoDB instances. This ensures that the DTS servers can connect to the two ApsaraDB for MongoDB instances. After the migration is complete, you can remove these IP addresses from the whitelists if you no longer need them. For more information, see [#unique_96](#).

7. Configure migration types and migration objects.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p> 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.</p> </div>

Parameter	Description
Available	<ul style="list-style-type: none"> In the Available section, select the objects you want to migrate and then click the  icon to move them to the Selected section. <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p> Note: Data in the admin database cannot be migrated even if this database is selected.</p> </div> <ul style="list-style-type: none"> 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. In the lower-right corner, click **Precheck**.



Note:

- A precheck is performed before the migration task starts. The migration task starts only after the precheck succeeds.
- If the precheck fails, click the  icon for each failed check item to view their details. Perform a precheck again after the failures are fixed.

9. After the precheck succeeds, click **Next**.

10. On the **Confirm Settings** page, set **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

Do not manually end a migration task. If you do so, the system may fail to migrate all data of the database. Wait until the migration task is complete.

- Incremental data migration

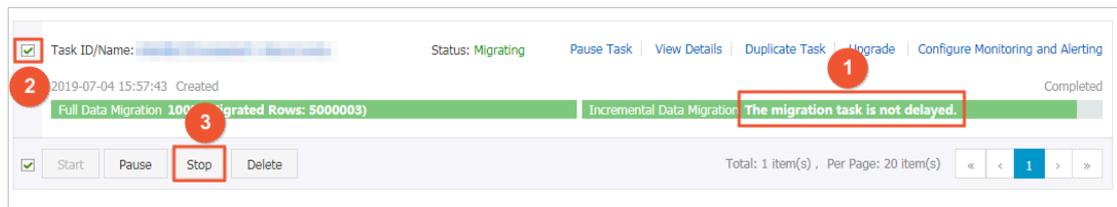
An incremental data migration task does not automatically end. You need to manually end the task.



Note:

Select an appropriate point in time to manually end a migration task. For example, you can end the migration task during off-peak hours or before you switch over your business to the destination ApsaraDB for MongoDB instance.

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



12. Switch over your business to the destination ApsaraDB for MongoDB instance.

8.10 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 [#unique_84](#).



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 [#unique_97](#), [#unique_98](#), or [#unique_99](#).

**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 [#unique_79](#).

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

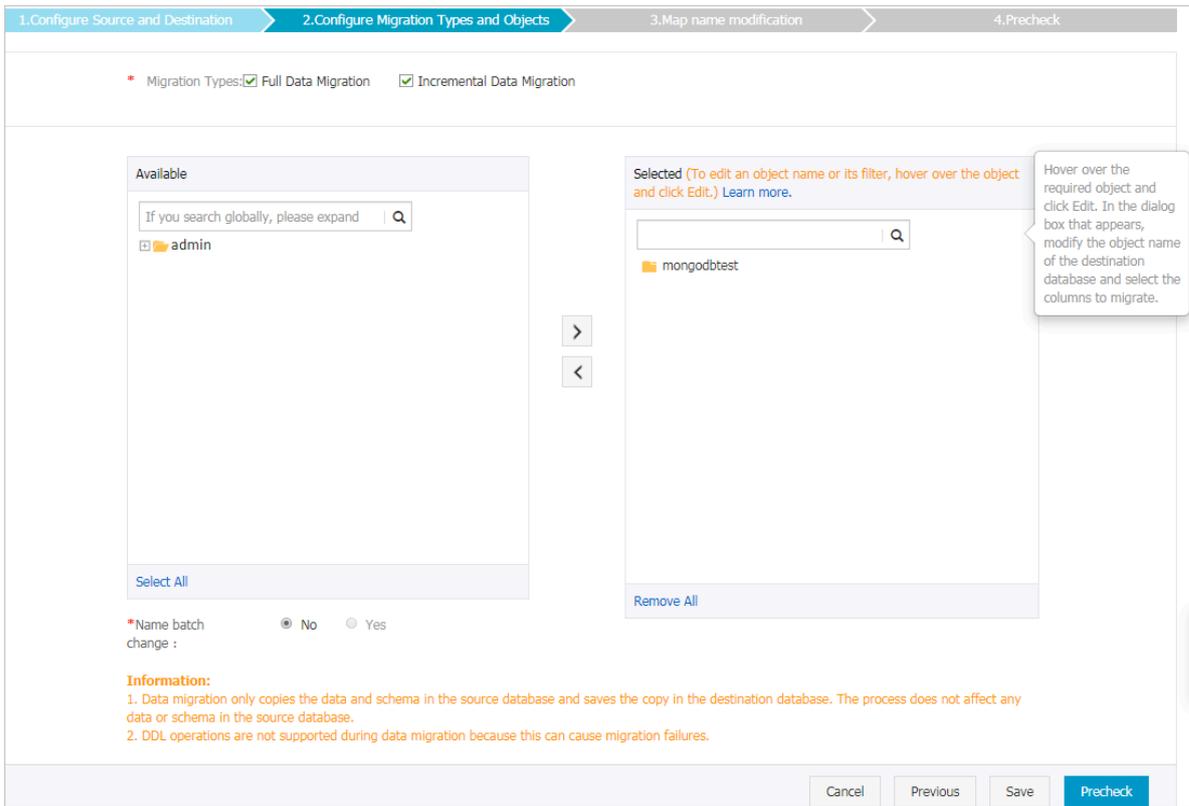
1. Configure both the source and destination databases.

The screenshot shows the '1. Configure Source and Destination' step in the DTS console. It features a progress bar at the top with four steps: 1. Configure Source and Destination (active), 2. Configure Migration Types and Objects, 3. Map name modification, and 4. Precheck. Below the progress bar, there is a 'Task Name' input field. The 'Source Database' section includes: Instance Type (ApsaraDB for MongoDB), Instance Region (China (Qingdao)), MongoDB Instance ID (dds-bp-...), Database Name (admin), Database Account, and Database Password. A 'Test Connectivity' button is located to the right of the password field. The 'Destination Database' section includes: Instance Type (MongoDB Instance), Instance Region (China (Hangzhou)), MongoDB Instance ID (dds-bj-...), Database Name (admin), Database Account, and Database Password. A 'Test Connectivity' button is also present here. At the bottom right, there are three buttons: 'Cancel', 'Assess Data Migration to Cloud', and 'Set Whitelist and Next'.

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify your own task name that helps identify the task. Task names do not need to be unique.
Source Database	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 Name	Enter the name of the authentication database. It is the database where the database account is created. <div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;">  Note: If the database account is root, enter admin. </div>

Section	Parameter	Description
	Database 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 .
	Database Password	Enter the password of the 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 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.
Destination Database	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.
	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 .
	Database 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.

2. Configure migration types and migration objects.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> 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. </div>

Parameter	Description
Available	<ul style="list-style-type: none">In the Available section, select the objects you want to migrate and then click the  icon to move them to the Selected section. <div data-bbox="443 412 1433 568"> Note: Data in the admin database cannot be migrated even if this database is selected.</div> <ul style="list-style-type: none">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.

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 [#unique_100](#).
- If the source instance uses subscription billing, you cannot release it.

9 Migrate data across Alibaba Cloud accounts

9.1 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](#).

Background information

- 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%.
- If your ApsaraDB RDS for MySQL database does not have primary key or unique constraints and each field in the database has duplicate values, the data migrated to the destination database may be duplicated.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the 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 [#unique_73](#).



Note:

You can use these three migration types together to migrate data without service interruptions.

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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

SQL operations that can be synchronized during incremental data migration

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul style="list-style-type: none"> 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

Database	Required permission
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 [#unique_72](#).

**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 **Migration Tasks** the 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**.

1.Configure Source and Destination > 2.Configure Migration Types and Objects > 3.Advanced Settings > 4.Precheck

* Task Name:

Source Database

* Instance Type: [DTS support type](#)

* Instance Region: [Get IP Address Segment of DTS](#)

* Database Type:

* Hostname or IP Address:

* Port Number:

* Database Account:

* Database Password:

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

* Task Name:

Source Database

* Instance Type: [DTS support type](#)

* Instance Region: [Guide](#)

* Apsara Stack Tenant Account ID:

* Role Name: [Authorize Role Across Accounts](#)

* Peer VPC: [Proprietary network of the current login account](#)

* Database Type:

* IP Address:

* Port Number:

* Database Account:

* Database Password: ✔ Passed

Destination Database

* Instance Type:

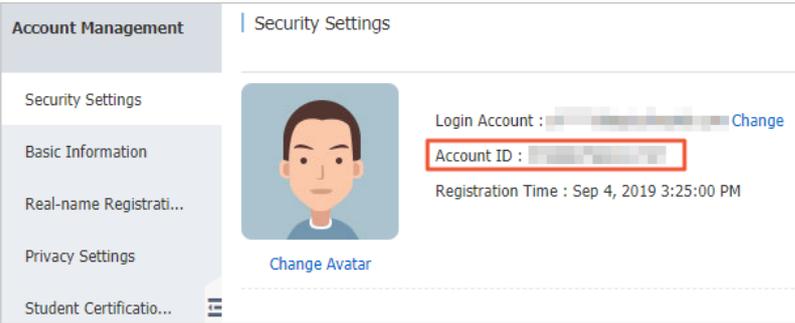
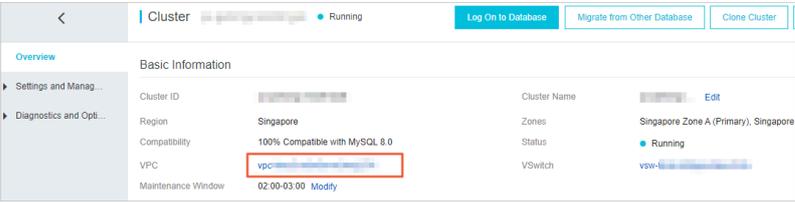
* Instance Region:

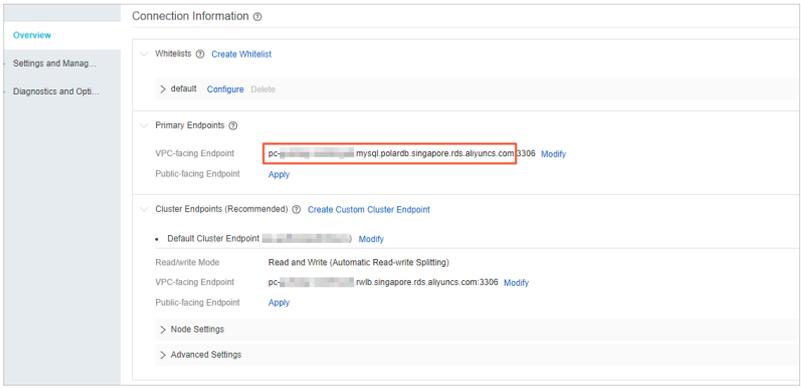
* PolarDB Instance ID:

* Database Account:

* Database Password: ✔ Passed

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	Select User-Created Database Connected over Express Connect, VPN Gateway, or Smart Access Gateway .
	Instance Region	Select the region where the source PolarDB cluster resides.

Section	Parameter	Description
	Alibaba Cloud Account ID	<p>Enter the ID of the Alibaba Cloud account to which the source PolarDB cluster belongs.</p> <div data-bbox="646 369 1433 607">  <p>Note: To obtain the ID of the Alibaba Cloud account to which the source PolarDB cluster belongs, you must log on to the Account Management console by using this account. The account ID is displayed on the Security Settings page.</p> </div> <div data-bbox="638 629 1433 952">  </div>
	Role Name	<p>Enter the name of the RAM role that you created earlier in Preparations.</p>
	Peer VPC	<p>Select the ID of the VPC where the source PolarDB cluster resides.</p> <p>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.</p> <div data-bbox="638 1473 1433 1675">  </div>
	Database Type	<p>Select MySQL.</p>

Section	Parameter	Description
	IP Address	<p>Enter the private IP address of the source PolarDB cluster. In this example, enter 172.16.20.20.</p> <p>You can obtain the private IP address by pinging the VPC-facing endpoint of the source PolarDB cluster.</p> 
	Port Number	Enter the service port number of the source PolarDB cluster. The default port number is 3306 .
	Database 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 Password	<p>Enter the password of the source database account.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <p>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.</p> </div>
Destination 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.

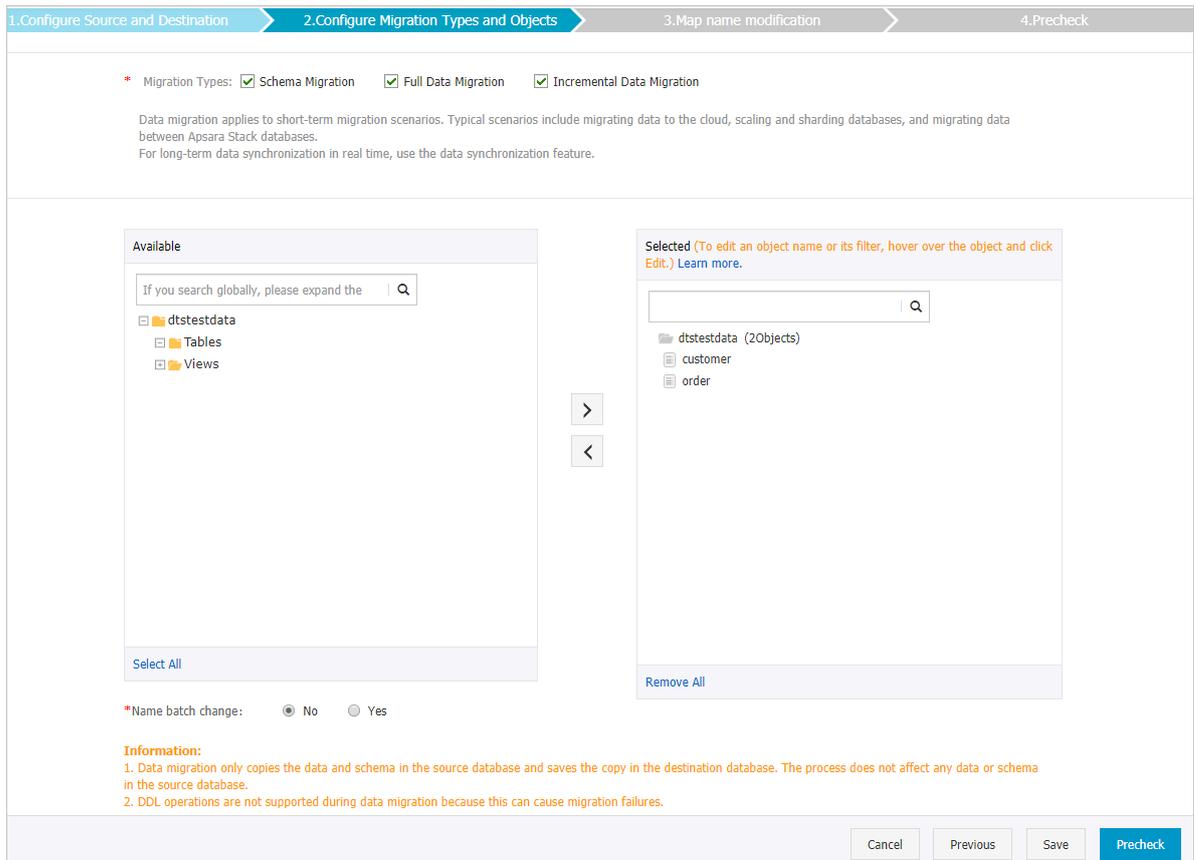
Section	Parameter	Description
	Database Account	Enter the database account of the destination PolarDB cluster. For more information about the permissions that are required for the account, see Permissions required for database accounts .
	Database Password	Enter the password of the destination database account.  Note: After you specify the destination database parameters, click Test Connectivity next to Database Password to verify whether the 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.

- Click **Set Whitelist and Next** in the lower-right corner of the page.

**Note:**

The IP addresses of DTS servers are added to the whitelist of the source and destination ApsaraDB for PolarDB clusters. This makes sure that the DTS server can connect to the clusters.

8. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> To perform only full data migration, select Schema Migration and Full Data Migration. If you want to migrate data without business disruptions, select Schema Migration, Full Data Migration, and Incremental Data Migration. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

9. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

10. After the precheck is passed, click **Next**.

11. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

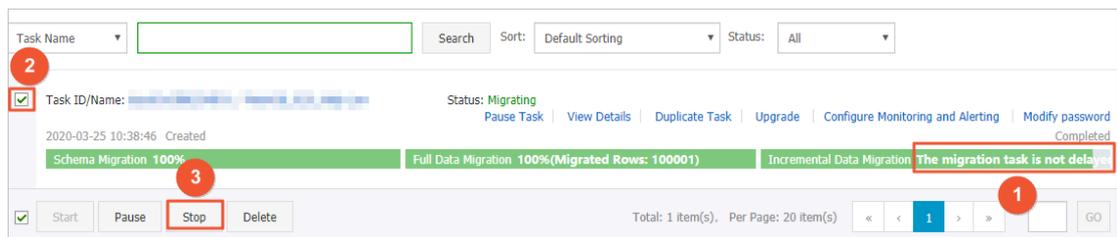
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- When 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 complete, you must delete the database accounts to ensure security.

9.2 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 [#unique_84](#).



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 [#unique_97](#), [#unique_98](#), or [#unique_99](#).



Note:

The storage capacity of the destination instance must be greater than the occupied storage space of the source instance.

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 [#unique_79](#).

Migration types

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

Data source	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

**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 [#unique_102](#).
3. Add the Classless Inter-Domain Routing (CIDR) blocks of DTS servers to a whitelist of the source instance. For more information, see [#unique_103](#).

**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. Configure both the source and destination databases.

The screenshot shows the configuration interface for DTS, divided into four steps: 1. Configure Source and Destination Databases, 2. Configure Migration Types and Objects, 3. Map name modification, and 4. Precheck. The 'Task Name' field is at the top. Below are two sections: 'Source Database' and 'Destination Database'. Each section contains fields for Instance Type, Instance Region, Database Type, Hostname or IP Address, Port Number, Database Name, Database Account, and Database Password. There are 'Test Connectivity' buttons for both. At the bottom right, there are buttons for 'Cancel', 'Assess Data Migration to Cloud', and 'Set Whitelist and Next'.

Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify your own task name that helps identify the task. Task names do not need to be unique.
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-1udxxxxxx-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.

Section	Parameter	Description
	Database 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 .
	Database Password	Enter the password of the 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 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.
Destination Database	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.
	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 .
	Database 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.

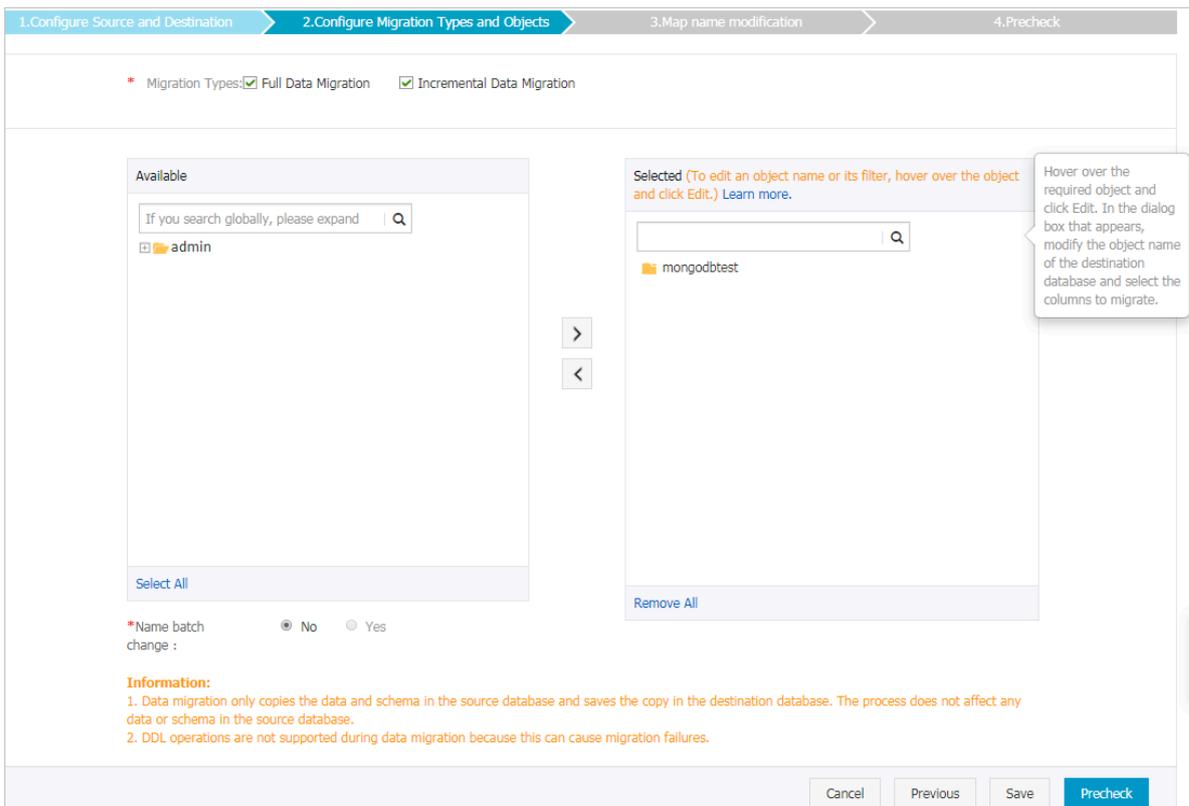
3. In the lower-right corner, click **Set Whitelist and Next**.



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](#).

4. Configure migration types and migration objects.



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <ul style="list-style-type: none"> 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. </div>

Parameter	Description
Available	<ul style="list-style-type: none"> In the Available section, select the objects you want to migrate and then click the  icon to move them to the Selected section. <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p> Note: Data in the admin database cannot be migrated even if this database is selected.</p> </div> <ul style="list-style-type: none"> 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.

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 [#unique_100](#).
- If the source instance uses subscription billing, you cannot release it.

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](#).

9.3 Migrate data between ApsaraDB for RDS instances created by different Alibaba Cloud accounts

This topic describes how to migrate Data from an Data Transmission Service instance to another Alibaba cloud account.

Prerequisites

The storage space of the destination instance must be larger than the occupied storage space in use of the source instance.

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 #unique_51 .
Incremental data migration	Charged. For more information, see #unique_51 .	

Permissions required for database accounts

Instance Type	Schema migration	Full data migration	Incremental data migration
Source ApsaraDB for RDS instance	The read/write permissions	The read/write permissions	The read/write permissions
Destination ApsaraDB for RDS instance	The read/write permissions	The read/write permissions	The read/write permissions

Preparations

Configure RAM authorization in the Alibaba cloud account to which the source instance belongs, and configure RAM authorization for the Alibaba Cloud account to which the destination instance belongs as a trusted account. Allow data transmission service to access cloud resources of the Alibaba Cloud account to which the source instance belongs. For more information, see [#unique_104](#).

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

* Task Name:

Source Database

* Instance Type: DTS support type

* Apsara Stack Tenant Account ID of RDS Instance: Guide

* Role Name: Authorize Role Across Accounts

* Instance Region: Retrieve Region List

* RDS Instance ID: RDS Instances of Current Account

* Database Account:

* Database Password: Test Connectivity

Destination Database

* Instance Type:

* Instance Region:

* RDS Instance ID:

* Database Account:

* Database Password: Test Connectivity

* Encryption: Non-encrypted SSL-encrypted

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.

Parameter	Description
Source Instance Details	<p>a. Behind the RDS Instance ID selection box, click RDS Instances of Other Apsara Stack Accounts.</p> <p>b. Configure the source instance.</p> <ul style="list-style-type: none"> • Instance Type: Select RDS Instance. • Apsara Stack Tenant Account ID of RDS Instance: Enter the ID of the Alibaba Cloud account to which the source instance belongs. <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;">  Note: Log on to the Alibaba Cloud account of the source instance. Account Management page to get the cloud account ID. </div> <ul style="list-style-type: none"> • Role Name: Enter the role name configured for the Alibaba Cloud account to which the source instance belongs. For more information, see #unique_104. • Instance Region: select the region where the source RDS instance is located. • RDS Instance ID: Select the ID of the source RDS instance. • Database Account: Enter the database account of the source instance. For more information about permissions required for the account, see Permissions required for database accounts. • Database Password: enter the password of the database account
Destination instance information	<ul style="list-style-type: none"> • Instance Type: Select RDS Instance. • Instance Region: select the region where the destination RDS instance is located. • 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. • Database Password: enter the password of the database account. • Connection mode: Select as needed Non-encrypted or SSL-encrypted in this case, Select Non-encrypted. <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;">  Note: 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 </div>

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



Note:

The IP addresses of DTS servers are automatically added to the whitelists of the source and destination RDS instances. This ensures that DTS servers can connect to the RDS instances.

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

Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>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.</p> </div>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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.

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



Note:

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

- Incremental data migration

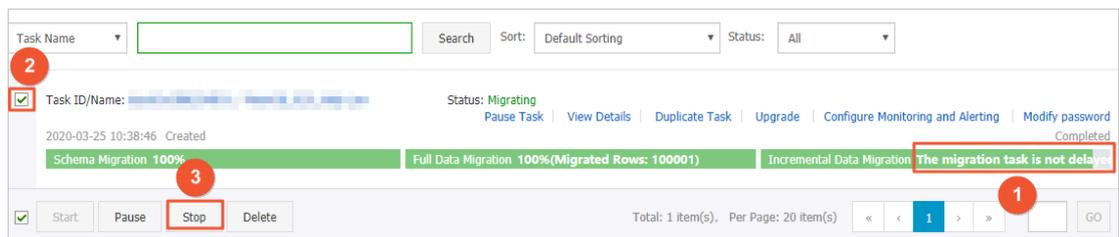
An incremental data migration task does not automatically stop. You must manually stop the migration task.



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**, manually stop the migration task.



10 Migration from Alibaba Cloud to a user-created database

10.1 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 UNIQUE 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.

Background information

- 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%.
- If your ApsaraDB RDS for MySQL database does not have primary key or unique constraints and each field in the database has duplicate values, the data migrated to the destination database may be duplicated.

- Concurrent insertions are performed during full data migration. This results in table fragmentation in the destination instance. After a full data migration task is completed, the tablespace of the destination instance is larger than that of the source instance.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the 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 of the following objects: tables, views, triggers, stored procedures, and stored functions.



Note:

During schema migration, the DEFINER mode of views, stored procedures, and stored functions is shifted to the INVOKER mode.

- The information of the source database account cannot be migrated. If you need to use views, stored procedures, and stored functions, you must grant read and write permissions to the destination database account.

Migration types

DTS supports schema migration, full data migration, and incremental data migration. For more information, see [#unique_73](#).



Note:

You can use these three migration types together to migrate data without service interruptions.

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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

SQL operations that can be synchronized during incremental data migration

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul style="list-style-type: none"> 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
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](#)
- [#unique_66](#) for a user-created MySQL database

Procedure

1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the 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.

1. Configure Source and Destination Databases
2. Configure Migration Types and Objects
3. Map name modification
4. Precheck

* Task Name:

Source Database

* Instance Type: DTS support type

* Instance Region:

* RDS Instance ID: RDS Instances of Other Apsara Stack Accounts

* Database Account:

* Database Password: ✔ Passed

* Encryption: Non-encrypted SSL-encrypted

Destination Database

* Instance Type:

* Instance Region:

* ECS Instance ID:

* Database Type:

* Port Number:

* Database Account:

* Database Password: ✔ Passed

* Encryption: Non-encrypted SSL-encrypted

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 Database	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.
	Database 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 .

Section	Parameter	Description
	Database Password	<p>Enter the password for the database account.</p> <p> Note: 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.</p>
	Encryption	<p>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.</p> <p> Note: The Encryption parameter is available only in mainland China and Hong Kong(China).</p>
Destination Database	Instance Type	Select User-Created Database in ECS Instance .
	Instance Region	Select the region where the ECS instance resides.
	ECS Instance ID	Select the ID of the ECS instance that is connected to the user-created MySQL database.
	Database Type	Select MySQL .
	Port Number	Enter the service port number of the user-created MySQL database. In this example, enter 3306 .
	Database 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 .

Section	Parameter	Description
	Database Password	<p>Enter the password for the account of the user-created MySQL database.</p> <div style="background-color: #f0f0f0; padding: 5px;">  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. </div>
	Encryption	Select Non-encrypted or SSL-encrypted . In this example, Non-encrypted is selected.

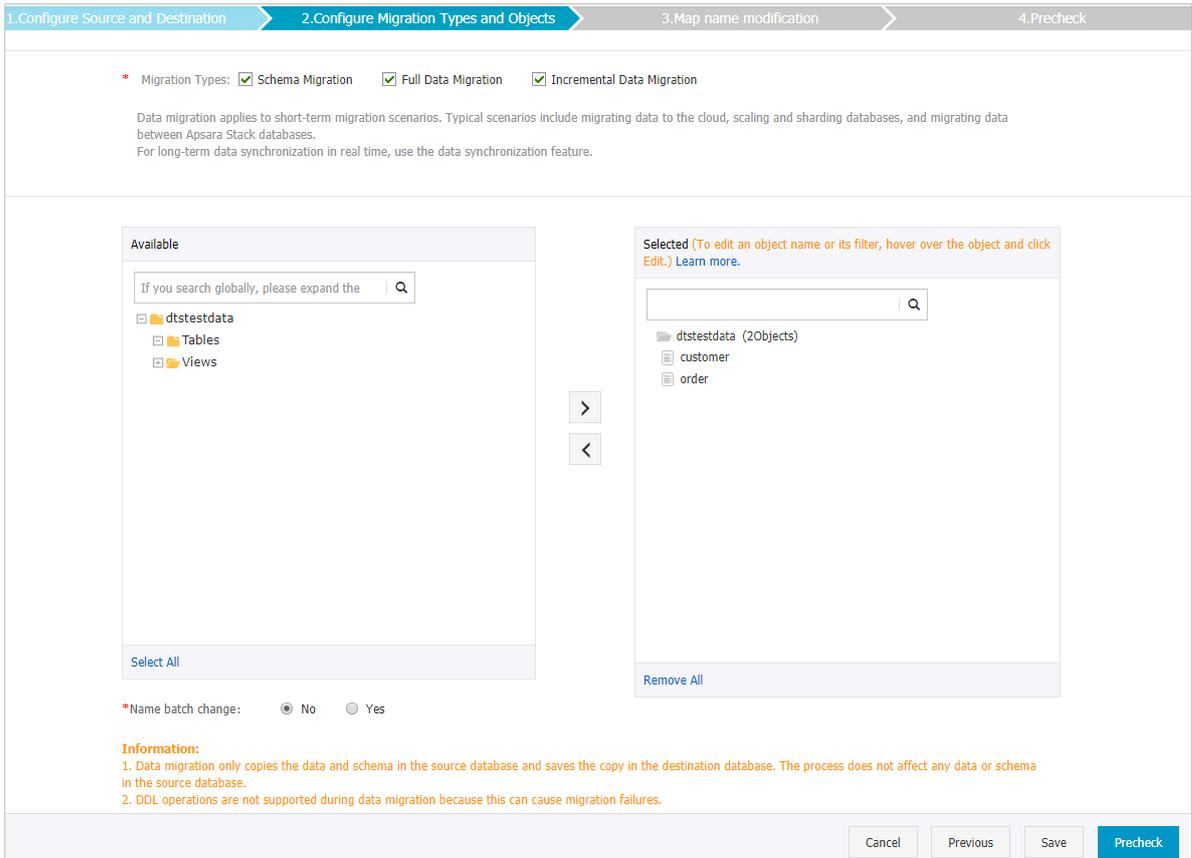
6. In the lower-right corner of the page, click **Set Whitelist and Next**.



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.



Item	Description
Migration types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

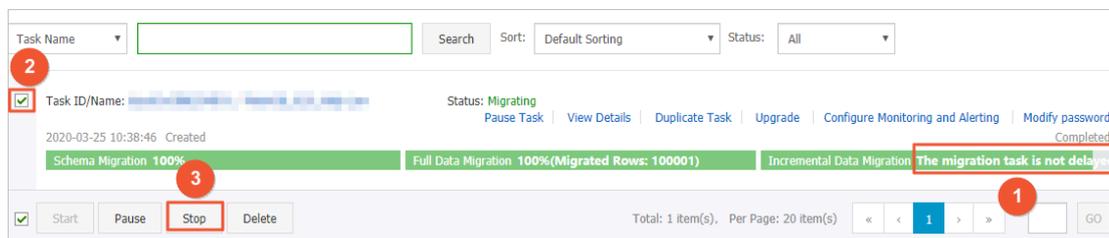
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



10.2 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](#).

Background information

- 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%.
- If your ApsaraDB RDS for MySQL database does not have primary key or unique constraints and each field in the database has duplicate values, the data migrated to the destination database may be duplicated.
- Concurrent insertions are performed during full data migration. This results in table fragmentation in the destination instance. After a full data migration task is completed, the tablespace of the destination instance is larger than that of the source instance.
- If a data migration task fails, DTS attempts to resume the task. In this case, before you switch your workloads to the destination database, you must stop or release the 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 of the following objects: tables, views, triggers, stored procedures, and stored functions.



Note:

During schema migration, the DEFINER mode of views, stored procedures, and stored functions is shifted to the INVOKER mode.

- The information of the source database account cannot be migrated. If you need to use views, stored procedures, and stored functions, you must grant read and write permissions to the destination database account.

Migration types

DTS supports schema migration, full data migration, and incremental data migration. For more information, see [#unique_73](#).



Note:

You can use these three migration types together to migrate data without service interruptions.

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 incurs fees. For more information, see #unique_51 .
Incremental data migration	Billed. For more information, see #unique_51 .	

SQL operations that can be synchronized during incremental data migration

Operation type	SQL statements
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	<ul style="list-style-type: none"> 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

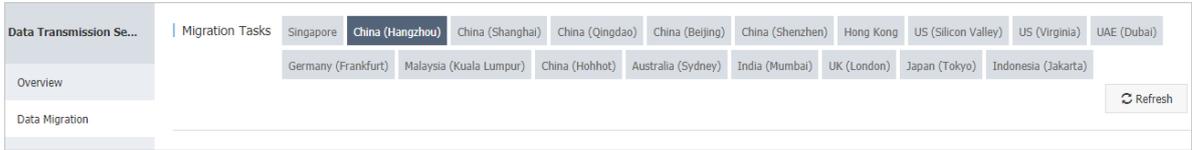
Database	Permission
Apsara PolarDB for MySQL cluster	The read permission for the objects to be migrated
User-created MySQL database	The read/write permissions for the objects to be migrated

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

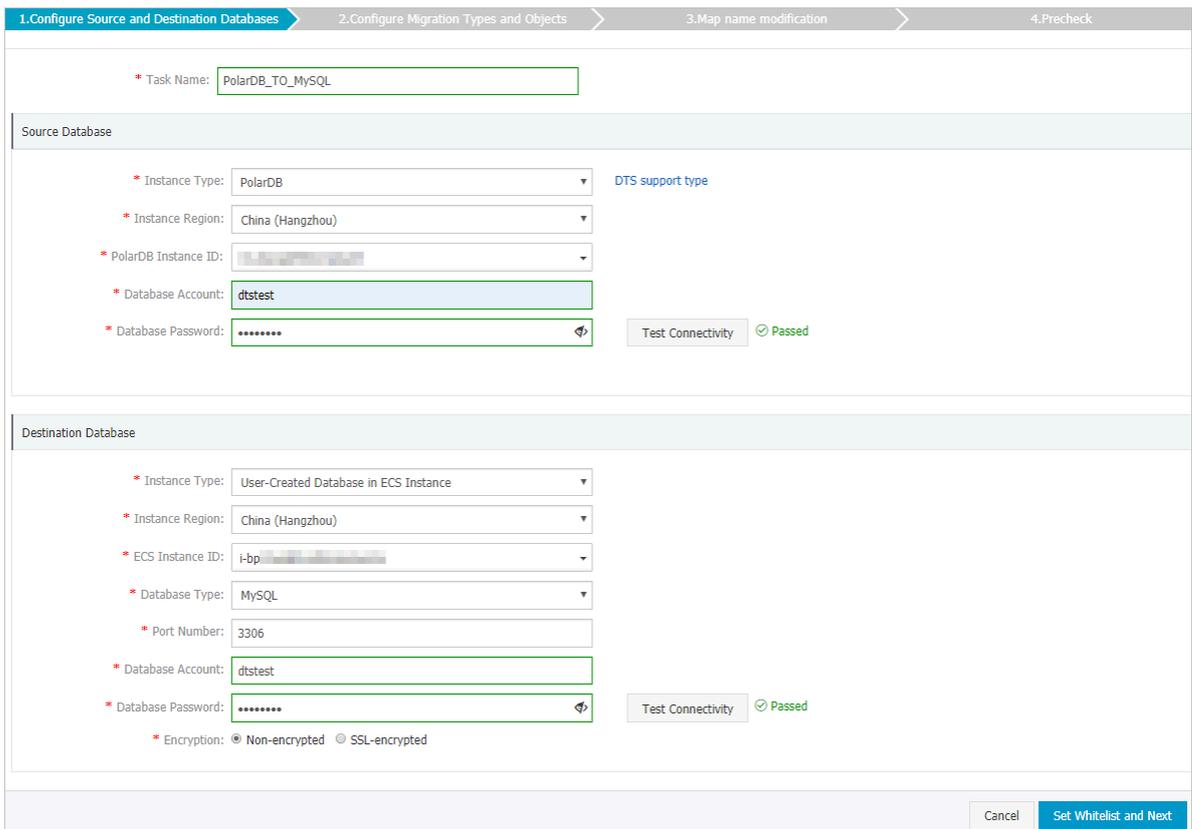
1. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.

3. At the top of **Migration Tasks** the 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.
Source Database	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
	Database 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 .
	Database Password	Enter the password for the database account.  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.
Destination Database	Instance Type	Select User-Created Database in ECS Instance .
	Instance Region	Select the region where the ECS instance resides.
	ECS Instance ID	Select the ID of the ECS instance that is connected to the user-created MySQL database.
	Database Type	Select MySQL .
	Port Number	Enter the service port number of the user-created MySQL database. In this example, enter 3306 .
	Database 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 .
	Database Password	Enter the password for the database account.  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.

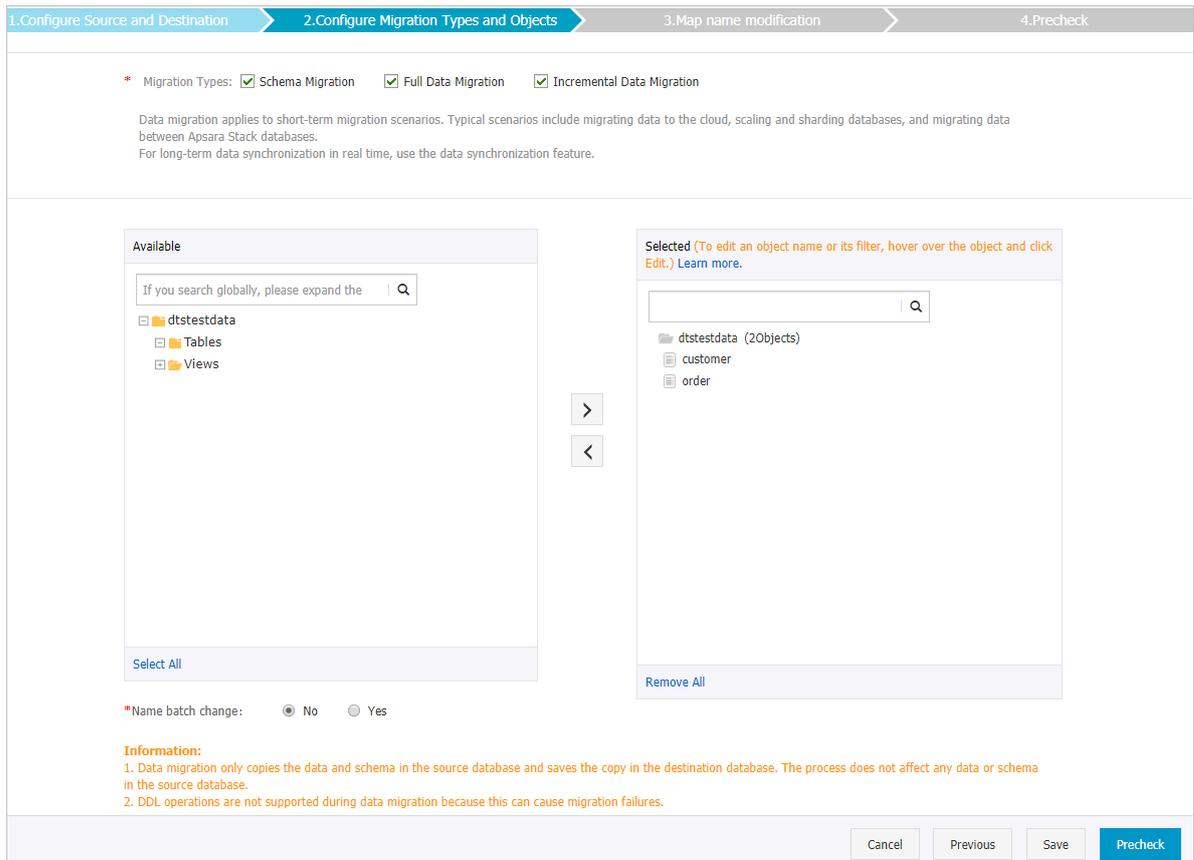
Section	Parameter	Description
	Encryption	Select Non-encrypted or SSL-encrypted . In this example, Non-encrypted is selected.

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. Configure migration types and objects.



Item	Description
Migration types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> 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.</p> </div>

Item	Description
Objects to be migrated	<p>Select the objects to be migrated in the Available section and click  icon to move them to the Selected section.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note:</p> <ul style="list-style-type: none"> • 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. </div>

8. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

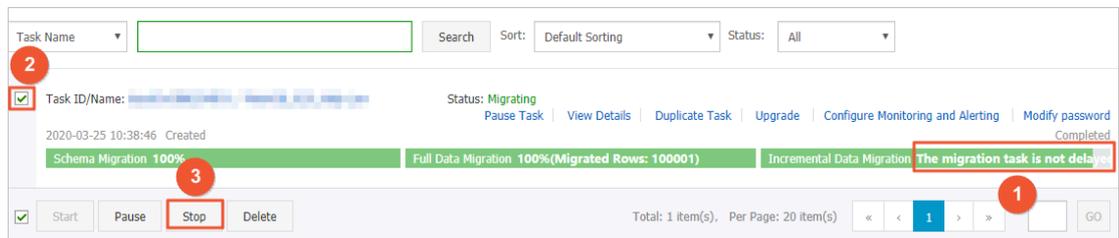
An incremental data migration task does not automatically end. You must manually end the migration task.



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- a. When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- b. When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



11 Migrate data between user-created databases

11.1 Migrate data between user-created 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 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 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	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.
Incremental data migration	Charged. For more information, see #unique_51 .	For more information, see #unique_51 .

Migration types

Migration type	Description
Schema migration	<p>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.</p> <div style="background-color: #f0f0f0; padding: 5px;">  Note: If an object contains triggers, the data between the source and destination databases will become inconsistent. </div>
Full data migration	<p>DTS migrates historical data of the required objects from the source Oracle database to the destination Oracle database.</p> <div style="background-color: #f0f0f0; padding: 5px;">  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. </div>

Migration type	Description
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.

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

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. Log on to the [DTS console](#).
2. In the left-side navigation pane, click **Data Migration**.
3. At the top of **Migration Tasks** the 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.

* Task Name:

Source Database

* Instance Type: DTS support type

* Instance Region:

* ECS Instance ID:

* Database Type:

* Port Number:

* Instance Type: Non-RAC Instance RAC Instance

* SID:

* Database Account:

* Database Password: Passed

Destination Database

* Instance Type:

* Instance Region:

* ECS Instance ID:

* Database Type:

* Port Number:

* Instance Type: Non-RAC Instance RAC Instance

* SID:

* Database Account:

* Database Password: Passed

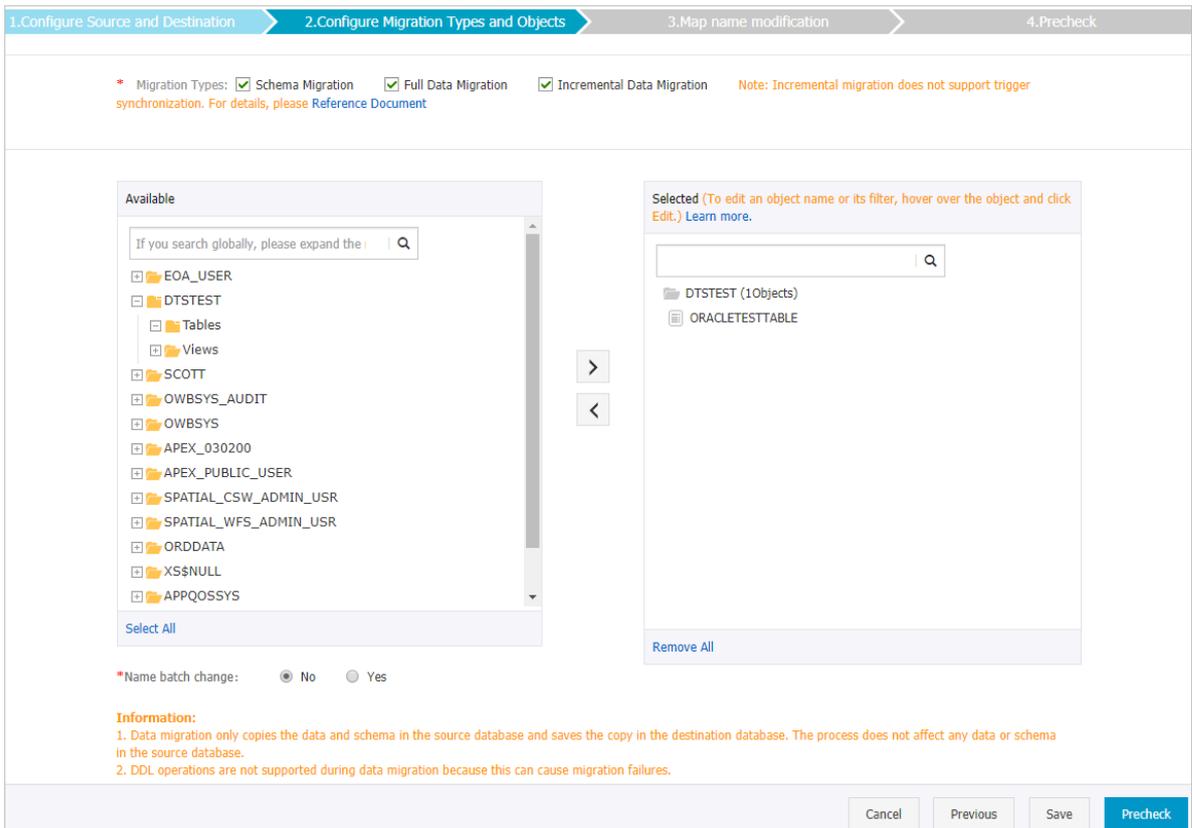
Section	Parameter	Description
N/A	Task Name	DTS automatically generates a task name. We recommend that you specify an informative name for easy identification . You do not need to use a unique task name.
Source Database	Instance Type	Select User-Created Database in ECS Instance . <div style="border: 1px solid #ccc; padding: 5px; background-color: #f0f0f0;"> <p> Note: If you select other instance types, you must prepare the environments that are required for the source database. For more information, see #unique_67.</p> </div>
	Instance 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 .

Section	Parameter	Description
	Port Number	Enter the service port number of the source Oracle database.
	Instance Type	<ul style="list-style-type: none"> • 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.
	Database 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;">  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. </div>
Destination Database	Instance Type	Select User-Created Database in ECS Instance .
	Instance Region	Select the region of the ECS instance on which the destination Oracle database is deployed.
	ECS Instance ID	Select the ID of the ECS instance on which the destination Oracle database is deployed.
	Database Type	Select Oracle .
	Port Number	Enter the service port number of the destination Oracle database.
	Instance Type	<ul style="list-style-type: none"> • 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.
	Database 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 .

Section	Parameter	Description
	Database Password	<p data-bbox="639 271 1406 300">Enter the password for the destination database account.</p> <div data-bbox="639 327 1433 689"> 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.</div>

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

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



Parameter	Description
Migration Types	<ul style="list-style-type: none"> 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. <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p> Note: If Incremental Data Migration is not selected, do not write data into the source Oracle database during full data migration. This ensures data consistency between the source and destination databases. For more information about migration types, see Migration types.</p> </div>

Parameter	Description
Objects	<p>Select objects from the Available section and click the  icon to move the objects to the Selected section.</p> <p> Note:</p> <ul style="list-style-type: none"> 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. Click **Precheck** on the lower right of the page.



Note:

- A precheck is performed for a data migration task. A data migration task can be started only if it passes the precheck.
- If the precheck fails, click  icon corresponding to each failed item to view the details. Fix the problems as instructed and run the precheck again.

9. After the precheck is passed, click **Next**.

10. On the **Confirm Settings** dialog box that appears, specify **Channel Specification** and select the **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

Do not manually stop a migration task. Otherwise, data migrated to the destination database will be incomplete. Wait until the data migration task stops when it is complete.

- Schema migration, full data migration, and incremental data migration

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



Note:

Select an appropriate time point to manually end the migration task. For example, you can end the migration task during off-peak hours or before you switch your workloads to the destination cluster.

- When the task progress bar switches to **Incremental Data Migration** and the message **The migration task is not delayed** appears, stop writing new data to the source database for a few minutes. Then, the progress bar will show the latency of the **incremental data migration**.
- When the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.

The screenshot displays the Data Migration Service console interface. At the top, there is a search bar with a dropdown for 'Task Name' and a 'Search' button. To the right, there are dropdowns for 'Sort: Default Sorting' and 'Status: All'. Below this, a task is listed with a checked checkbox and 'Task ID/Name: [redacted]'. The status is 'Migrating'. Action links include 'Pause Task', 'View Details', 'Duplicate Task', 'Upgrade', 'Configure Monitoring and Alerting', and 'Modify password'. The task creation time is '2020-03-25 10:38:46 Created'. The progress bar shows three stages: 'Schema Migration 100%' (green), 'Full Data Migration 100%(Migrated Rows: 100001)' (green), and 'Incremental Data Migration' (green) with a red box around the text 'The migration task is not delayed'. At the bottom, there are 'Start', 'Pause', 'Stop', and 'Delete' buttons. The 'Stop' button is highlighted with a red box and a red circle with the number '3'. A red circle with the number '1' is positioned over the pagination controls, which show 'Total: 1 item(s), Per Page: 20 item(s)' and a page number '1'.