Alibaba Cloud

Data Transmission Service Quick Start

Document Version: 20220712

(-) Alibaba Cloud

Legal disclaimer

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

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
- 4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
- 5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- 6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

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

> Document Version: 20220712

Table of Contents

1.Purchase procedure	05
2.Manage a data synchronization task	09
3.Manage a data migration task	10
4.Manage a change tracking task	11
5.Configuration guide	14
5.1. Configure a data synchronization task	14
5.2. Configure a data migration task	17
5.3. Configure a data synchronization task	20
5.4. Create a data migration task	24
5.5. Create a change tracking task	27

1.Purchase procedure

This topic describes how to purchase a data migration instance, a data synchronization instance, and a change tracking instance in Data Transmission Service.

Prerequisites

- An Alibaba Cloud account is registered. For more information, see Sign up with Alibaba Cloud.
- Your account balance is sufficient if you want to create a pay-as-you-go instance.

Billing

For more information, see Pricing.

Purchase a data migration instance

Note You can also skip the purchase step and directly create a data migration task. For more information, see Create a data migration task.

- 1. Open purchase page of Data Transmission Service.
- 2. Select the billing method to Pay-As-You-Go.
- 3. Configure parameters for the data migration instance. The following table describes the parameters.

Section	Parameter	Description
	Function	Select Data Migration.
Pacie	Region	Select the region where the destination instance resides.
Basic	Specification	DTS provides several specifications of migration channels that have different performance. For more information, see Specifications of data migration instances.
Purchase Plan	Number	Set the number of data migration instances that you want to purchase at a time. The default value is 1, you can purchase a maximum number of 99 instances at a time.

4. Click **Buy Now**, and follow the instructions to complete the payment.

Purchase a data synchronization instance

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click **Data Synchronization**.
- 3. In the upper-right corner of the page, click Create Synchronization Task.
- 4. On the purchase page of Data Transmission Service, select a billing method.
 - Subscription: You need to first pay for the subscription when creating an instance. We
 recommend that you select this billing method for long-term use because it is more costeffective than the pay-as-you-go billing method. You can save more with longer subscription
 periods.

- Pay-as-you-go: You need to pay for a service after you use the service. You are billed on an hourly basis. We recommend that you select this billing method for short term use. You can save costs by deleting the instance when it is no longer required.
- 5. Configure parameters for the data synchronization instance. The following table describes the parameters.

Section	Parameter	Description
_	Function	Select Data Synchronization.
	Source Instance	Select the type of the source instance that is used for data synchronization.
	Source Region	Select the region where the source instance resides. You cannot change the region after purchase.
	Target Instance	Select the type of the destination instance that is used for data synchronization.
	Target Region	Select the region where the destination instance resides. You cannot change the region after purchase.
		Select One-Way Synchronization or Bidirectional Synchronization.
Basic	Synchroniz at io n Topology	 Note For more information about the synchronization topology, see Databases, initial synchronization types, and synchronization topologies. Bidirectional Synchronization is only available when you select MySQL for Source Instance and Target Instance.
	Specification	DTS provides several specifications of synchronization channels that have different performance. The performance is measured by the number of data records that can be synchronized. For more information, see Specifications of data synchronization instances. ? Note In the production environment, we recommend that you select small or a higher specification.
		If you select the subscription billing method, set the subscription
	Order Time	period. The subscription period can be one to nine months or one to three years. Note The configuration of the subscription period is only available when the billing method is subscription.
Purchase Plan		only available when the bining method is subscription.

Section	Parameter	Description
	Quantity	Set the number of data synchronization instances that you want to purchase at a time. The default value is 1. If the billing method is pay-as-you-go, you can purchase a maximum number of 99 instances at a time.

6. Click **Buy Now**, and follow the instructions to complete the payment.

Both of your source and destination instances must be located in mainland China or in the same country. (Note: If the source instance is located in Hong Kong, the destination instance must be located in Hong Kong.) If you want to place your instances in different countries, submit a ticket or contact the customer sales manager.

Next step

Configure a data synchronization task (旧控制台)

Purchase a change tracking instance

- 1. Log on to the DTS console.
- 2. In the left-side navigation pane, click Change Tracking.
- 3. In the upper-right corner of the page, click Create Change Tracking Task.
- 4. On the purchase page of Data Transmission Service, select a billing method.
 - Subscription: You need to first pay for the subscription when creating an instance. We
 recommend that you select this billing method for long-term use because it is more costeffective than the pay-as-you-go billing method. You can save more with longer subscription
 periods.
 - Pay-as-you-go: You need to pay for a service after you use the service. You are billed on an hourly basis. We recommend that you select this billing method for short term use. You can save costs by deleting the instance when it is no longer required.
- 5. Configure parameters for the change tracking instance. The following table describes the parameters.

Section	Parameter	Description
	Function	Select Data Subscription.
Basic	Region	Select the region where the source database resides. You cannot change the region after purchase.
C	Order Time	If you select the subscription billing method, set the subscription period. The subscription period can be one to nine months or one to three years.
		Note The configuration of subscription period is only available when the billing method is subscription.
Purchase Plan		

Section	Parameter	Description
	Quantity	Set the number of change tracking instances that you want to purchase at a time. The default value is 1. If the billing method is pay-as-you-go, you can purchase a maximum number of 99 instances at a time.

6. Click ${\bf Buy\ Now}$, and follow the instructions to complete the payment.

Next step

Create a change tracking task

2.Manage a data synchronization task

You can use Data Transmission Service (DTS) to synchronize data between two data sources in real time. This feature applies to multiple scenarios, such as active geo-redundancy, geo-disaster recovery, zone-disaster recovery, cross-border data synchronization, query load balancing, cloud business intelligence (BI) systems, and real-time data warehousing. This topic describes how to create, monitor, and manage a data synchronization task.

Background information

• Overview of data synchronization scenarios

This topic describes the database types, initial synchronization types, and synchronization topologies that are supported by DTS.

• Data type mappings for schema synchronization

The data types of different databases do not have one-to-one correspondence. This topic describes the data type mappings that DTS uses between the source and destination databases during initial schema synchronization. Based on the data type mappings, you can evaluate the impact of data type conversion on your business.

Features

For information about the features of a data synchronization task, Data synchronization.

Manage a data synchronization task

Step	Description
1. Prepare for data synchronization	Before you use DTS for data synchronization, prepare the environments that are required for the source and destination databases.
2. Configure a data synchronization task (旧控制台)	Create a data synchronization task by completing simple configurations in the DTS console.
3. (Optional) Manage the data synchronization task when it is running	 View the connection status and performance of data synchronization Monitor the data synchronization task on the dashboard that is provided by DTS. Add an object to the data synchronization task or Remove an object from the data synchronization task Add or remove objects from the data synchronization task.
4. Stop the data synchronization task	Stop the data synchronization task if you no longer need the task.

Examples

Overview of data synchronization scenarios

3. Manage a data migration task

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 how to create, monitor, and manage a data migration task.

Background information

Overview of data migration scenarios

This topic describes the database types, database versions, and migration types that are supported by DTS.

• Data type mappings between heterogeneous databases

The data types of different databases do not have one-to-one correspondence. This topic describes the data type mappings that DTS uses between the source and destination databases during schema migration. You can evaluate the impact of data type conversion on your business.

Features

For information about the features of a data migration task, see Data migration.

Manage a data migration task

Step	Description
1. Prepare for data migration	Before you use DTS for data migration, prepare the environments that are required for the source and destination databases.
2. Create a data migration task	Create a data migration task by completing simple configurations in the DTS console.
3. View the progress of the data migration task	View the progress of schema migration, full data migration, or incremental data migration in the DTS console.
4. Switch workloads to the destination database	Switch your workloads to the destination database. Before you switch the workloads, we recommend that you prepare a rollback solution to minimize the negative impacts of data migration on your business.
5. Stop the data migration task	After data migration and workload switchover is completed, stop the data migration task if you no longer need the task.

Examples

Overview of data migration scenarios

4. Manage a change tracking task

You can use DTS to retrieve incremental data from data sources in real time. This feature applies to multiple scenarios, such as cache update, business decoupling, asynchronous data processing, real-time synchronization of heterogeneous data, and real-time synchronization of complex extract, transform, and load (ETL) operations. This topic describes how to create, monitor, and manage a change tracking task.

Background information

• Overview of change tracking scenarios

This topic describes the databases and data types that are supported by the change tracking feature.

Features

For information about the features of a change tracking task, see Change tracking (new) and Change tracking (previous).

Manage a change tracking task

Step	Supported version of the change tracking feature	Description
1. Prepare for change tracking	New and previous	Before you use DTS for change tracking, prepare the environments that are required for the source and destination databases.
2. Create a change tracking task	New and previous	Create a change tracking task by completing simple configurations in the DTS console.

Step	Supported version of the change tracking feature	Description
3. Manage the change tracking task	New	 Create consumer groups After you configure a change tracking task, create one or more consumer groups to consume data. You can create multiple consumer groups for a change tracking task. Consumers in different consumer groups can track data changes from the same data source. Consumer groups allow you to reduce the cost for tracking data changes and improve the efficiency of data consumption. (Optional)Manage consumer groups Manage the consumer groups of the change tracking task. For example, you can modify passwords for consumer groups or delete consumer groups. (Optional)Modify the objects for change tracking Change the objects whose changes are tracked when the change tracking task is running. (Optional)View tracked data changes View the data changes that DTS retrieves from the data source.
	Previous	 (Optional) (Optional)Modify objects for change tracking Change the objects whose changes are tracked when the change tracking task is running. (Optional)Edit the consumption checkpoint Change the consumption checkpoint when the change tracking task is running. Downstream SDK clients will consume data that is generated after the checkpoint. View tracked data changes View the data changes that DTS retrieves from the data source.
	New	Use a Kafka client to consume tracked data
4. Consume tracked data.	Previous	Use SDK to consume data

Step	Supported version of the change tracking feature	Description
5. Stop the change tracking task	New and previous	 If you want to stop the change tracking task, delete the task. Pay-as-you-go: For information about how to stop a pay-as-you-go change tracking task, see Release pay-as-you-go instances. Subscription: A subscription change tracking task cannot be manually deleted. It will be automatically deleted seven days after the expiration date.
		Note You can also reset a change tracking task to stop the task. For more information, see Reset a change tracking task.

Examples

Overview of change tracking scenarios

5.Configuration guide5.1. Configure a data synchronization task

Data Transmission Service (DTS) provides an easy-to-use data synchronization feature. To configure a data synchronization task, you only need to perform simple operations in the DTS console.

Precautions

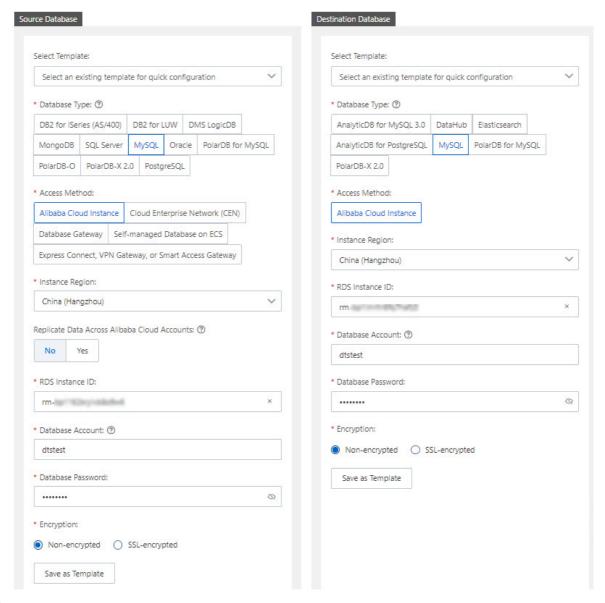
This topic provides a general procedure for configuring a data synchronization task. The procedure varies with the source and destination databases.

Before you begin

- If the source or destination database is a self-managed database, you must deploy the network environment for the database. For more information, see Preparation overview.
- Make sure that the source and destination instances are created and the database versions are supported by DTS.

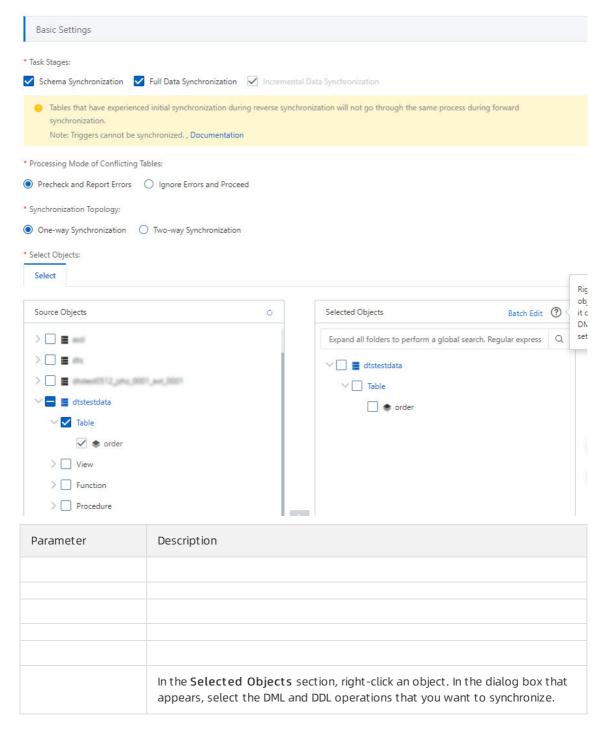
Procedure

- 1.
- 2.
- 3.
- 4.

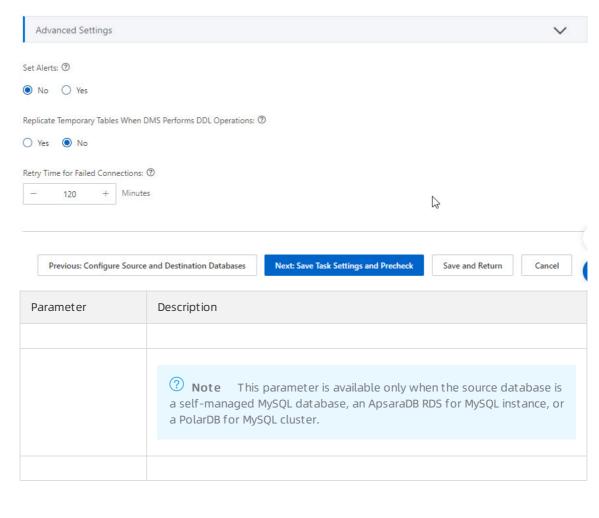


6.

7. ○ Basic Settings



Advanced Settings



8.

9.

10.

11.

12.

5.2. Configure a data migration task

Data Transmission Service (DTS) provides an easy-to-use data migration feature. To configure a data migration task, you only need to perform simple operations in the DTS console.

Precautions

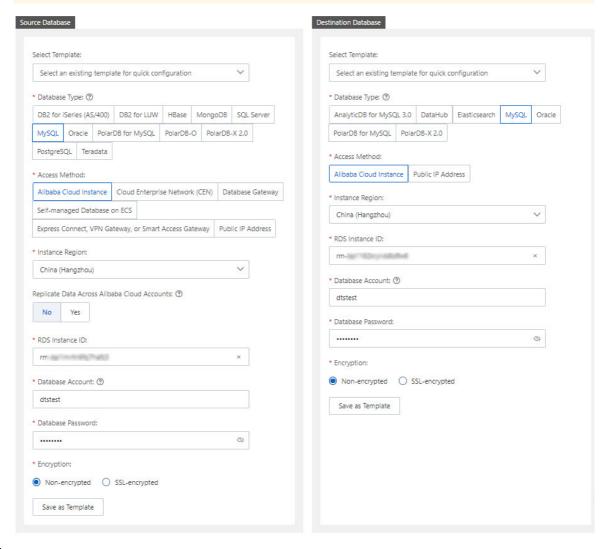
This topic provides a general procedure for configuring a data migration task. The procedure varies with the source and destination databases.

Before you begin

- If the source or destination database is a self-managed database, you must deploy the network environment for the database. For more information, see Preparation overview.
- Make sure that the source and destination instances are created and the database versions are supported by DTS.

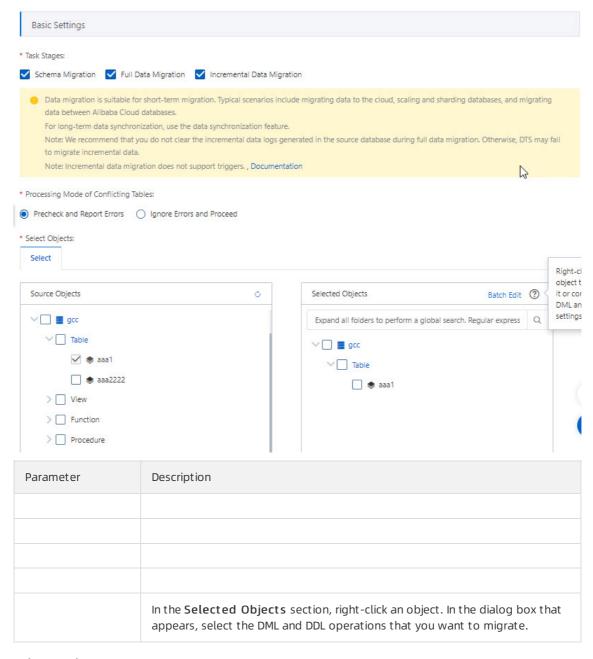
Procedure

- 1.
- 2.
- 3.
- 4.

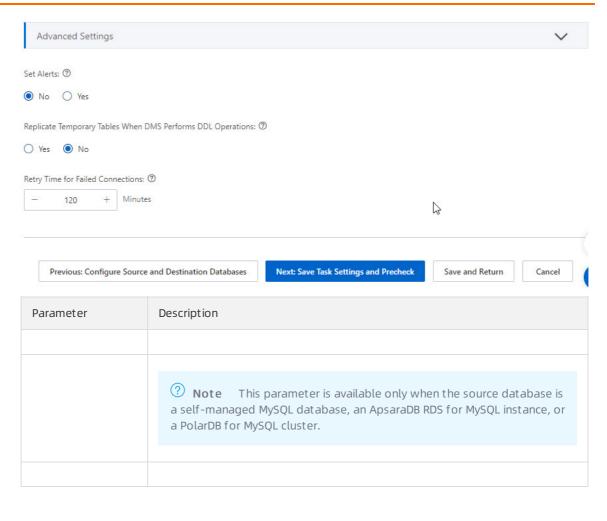


6.

7. ○ Basic Settings



Advanced Settings



8.

9.

10.

11.

12.

5.3. Configure a data synchronization task

Data Transmission Service (DTS) provides an easy-to-use data synchronization feature. To configure a data synchronization task, you only need to perform simple operations in the DTS console.

If you use the data synchronization feature for the first time, read the topic Manage a data synchronization task to learn how to create, monitor, and manage a data synchronization task.

Precautions

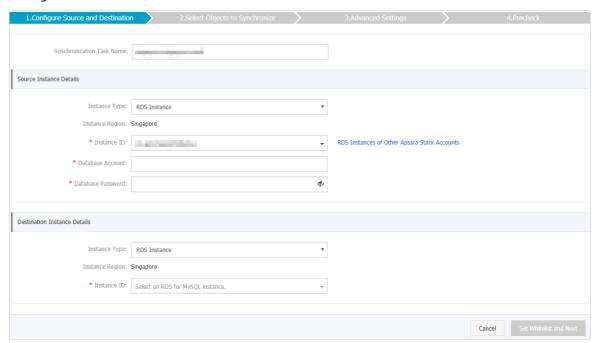
This topic provides a general procedure for configuring a data synchronization task. The procedure varies with the data sources. For more information about how to configure data synchronization tasks for specific data sources, see Overview of data synchronization scenarios.

Before you begin

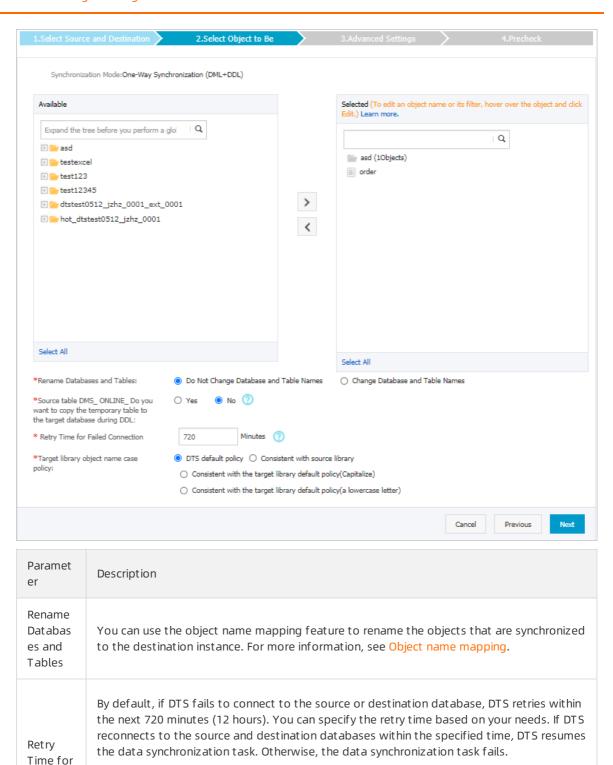
If the source database is a self-managed MySQL database or runs on a PolarDB for MySQL cluster, you must enable the binary logging feature for the database. For more information, see Enable binary logging and Create an account for a user-created MySQL database and configure binary logging.

Procedure

- 1. Purchase a data synchronization instance based on the database type and region of the source and destination instances. For more information, see Purchase procedure.
- 2. Log on to the DTS console.
- 3. In the left-side navigation pane, click **Data Synchronization**.
- 4. At the top of the **Synchronization Tasks** page, select the region where the data synchronization instance resides.
- 5. Find the data synchronization instance and click **Configure Synchronization Channel** in the Actions column.
- 6. Configure the source and destination instances.



- 7. In the lower-right corner of the page, click Set Whitelist and Next.
 - **? Note** If the source instance or destination instance is an Alibaba Cloud database instance, DTS automatically adds the CIDR blocks of DTS servers to the whitelist of the database instance. This ensures that DTS servers can connect to the database instance.
- 8. Select the synchronization policy and the objects to be synchronized. Click ${\bf Next}$.



Note When DTS retries a connection, you are charged for the DTS instance. We

recommend that you specify the retry time based on your business needs. You can

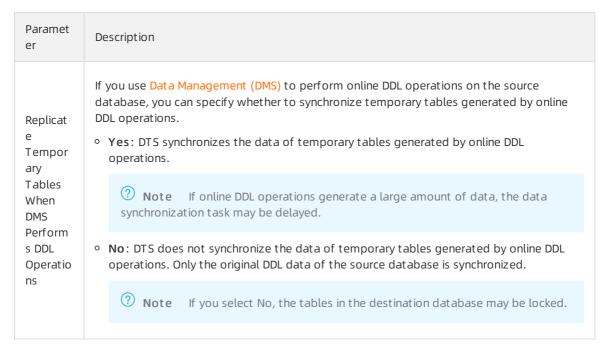
also release the DTS instance at your earliest opportunity after the source and

destination instances are released.

Failed

ons

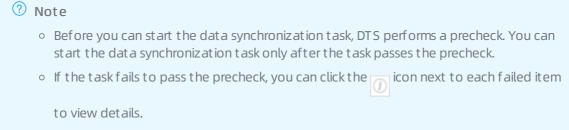
Connecti



9. Select the initial synchronization types.

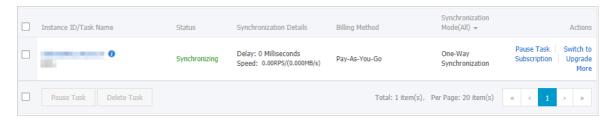


10. In the lower-right corner of the page, click **Precheck**.



- You can troubleshoot the issues based on the causes and run a precheck again.
- If you do not need to troubleshoot the issues, you can ignore failed items and run a precheck again.
- 11. Close the **Precheck** dialog box after the following message is displayed: **The precheck is passed.** Then, the data synchronization task starts.
- 12. Wait until initial synchronization is completed and the data synchronization task is in the **Synchronizing** state.

You can view the status of the data synchronization task on the **Synchronization Tasks** page.



References

- Overview of data synchronization scenarios
- Data type mappings for schema synchronization
- Synchronization topologies

5.4. Create a data migration task

Data Transmission Service (DTS) provides a simple and easy-to-use data migration feature. You only need to perform simple operations in the DTS console to configure the entire migration task.

Precautions

This topic only provides a general procedure for configuring a data migration task. However, the procedure will vary according to the data source. For more information about how to configure data migration tasks for specific data sources, see Overview of data migration scenarios.

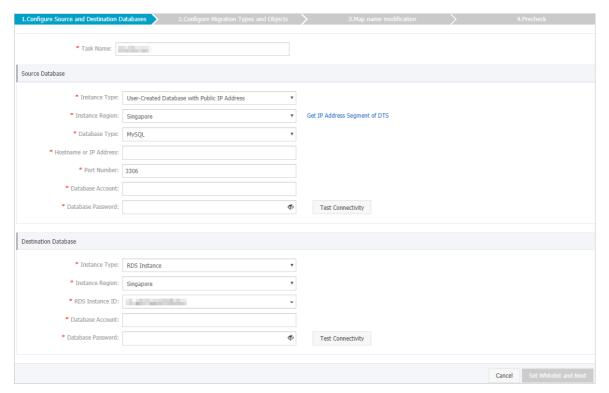
Prepare database accounts

When configuring a migration task, you must provide the accounts for accessing the source database and destination database.

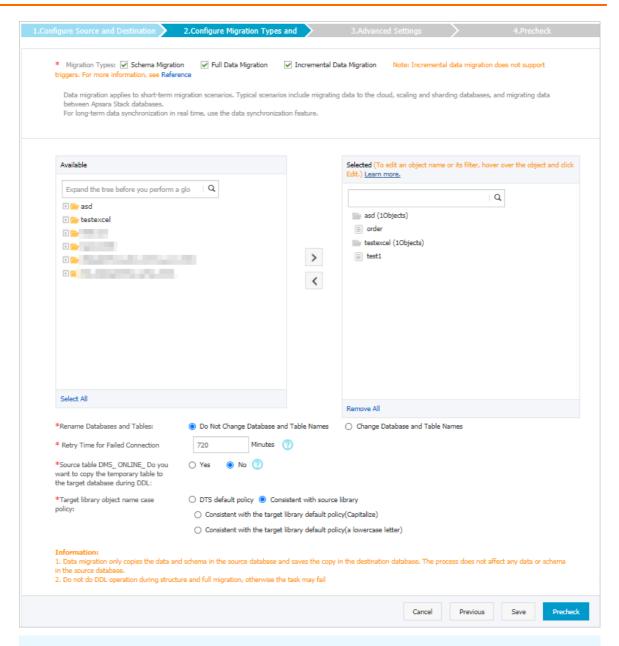
Procedure

24

- 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 database resides.
- 4. In the upper-right corner, click Create Migration Task.
- 5. Specify a task name and the connection details of the source database and destination database.



- 6. Click Set Whitelist and Next.
 - Note If the source database or destination database is in an Alibaba Cloud instance, the IP address of the DTS server is automatically added to the whitelist. This ensures that the DTS server can connect to the instance.
- 7. Specify migration types and objects to be migrated.



Note If the migrated object has different names in the source instance and destination instance, you can use the object name mapping function provided by DTS. For more information, see Object name mapping.

8. Click Precheck.

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



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

- 9. After the precheck is passed, click **Next**.
- 10. On the Confirm Settings page, specify Channel Specification, read the Data Transmission Service (Pay-As-You-Go) Service Terms, and then select the check box to agree to the terms.
- 11. Click **Buy and Start** to start the migration task.
 - o Full dat a migration

Do not manually stop a migration task. We recommend that you do not interrupt the migration task because the system may fail to migrate the full data of the database. Wait until the migration task stops when it is complete.

o Incremental data migration

An incremental data migration task does not automatically stop when the migration is complete. In this case, you need to manually stop the migration task.

- a. When the task progress bar displays Incremental Data Migration and The migration task is not delayed, stop writing to the source database for a few minutes. At this time, the progress bar will display the delay time of the incremental data migration.
- b. After the status of **incremental data migration** changes to **The migration task is not delayed**, manually stop the migration task.



5.5. Create a change tracking task

Data Transmission Service (DTS) provides a simple and easy-to-use change tracking feature. You only need to perform simple operations in the DTS console to configure the entire change tracking task. The change tracking feature allows you to subscribe to data changes in real time. It helps you achieve lightweight cache updates, business decoupling, asynchronous data processing, and real-time data synchronization through ETL operations.

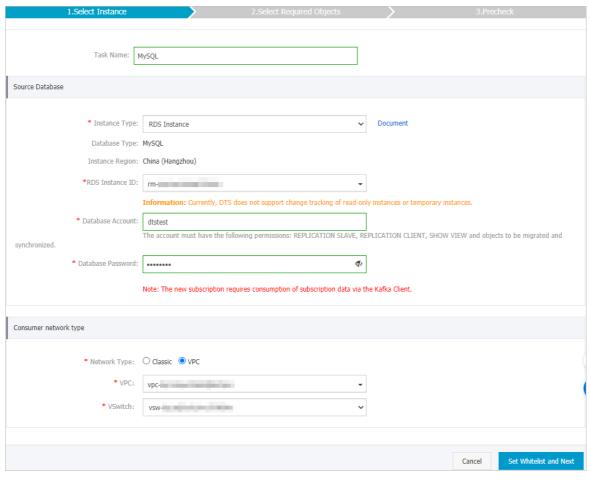
Precautions

This topic only provides a general procedure for configuring a change tracking channel. However, the procedure will vary according to the data source. For more information about how to configure change tracking channels for specific data sources, see Overview of change tracking scenarios.

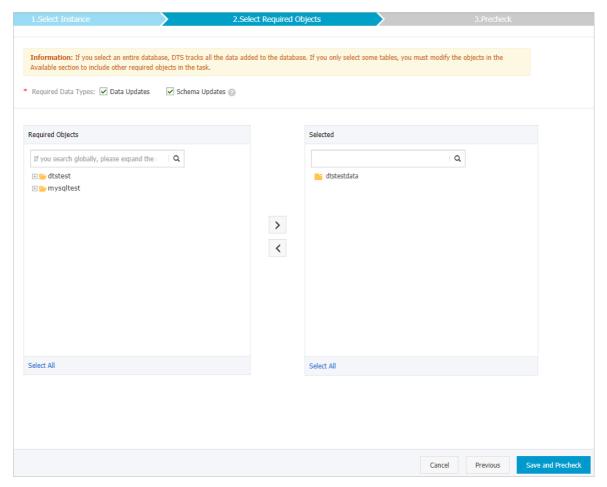
Procedure

- 1. Create a change tracking task by specifying the region of the source database that is used for data subscription. For more information, see Purchase a change tracking instance.
- 2. Log on to the DTS console.
- 3. In the left-side navigation pane, click Change Tracking.
- 4. At the top of the **Change Tracking Tasks** page, select the region to which your change tracking task belongs.
- 5. Find the created change tracking task and click **Configure Channel**.
- 6. Specify the connection details of the source database that is used for data subscription.

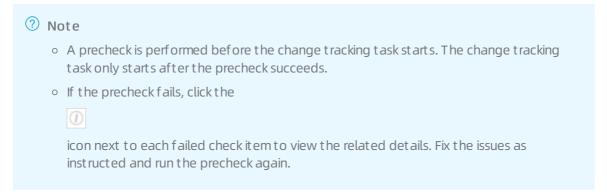
Note For more information about the previous and new version of the change tracking feature, see Change tracking (previous) and Change tracking (new).



- 7. Click Set Whitelist and Next.
 - Note If the source database is in an Alibaba Cloud instance, the IP address of the DTS server is automatically added to the whitelist for the instance. This ensures that the DTS server can connect to the instance.
- 8. Specify the data type and objects to which you want to subscribe.



9. Click Save and Precheck in the lower-right corner.



10. If the Precheck dialog box displays The precheck is passed, close the Precheck dialog box.

After the configuration is complete, the initial change tracking starts. This process requires about 1 minute to complete.

References

Change tracking version	Reference
Previous version	View tracked data changesSubscribe to data changes by using SDKs

Change tracking version	Reference
New version	 Create consumer groups View tracked data changes Use a Kafka client to consume tracked data