TDMQ for RocketMQ

Migration to Cloud

Product Documentation

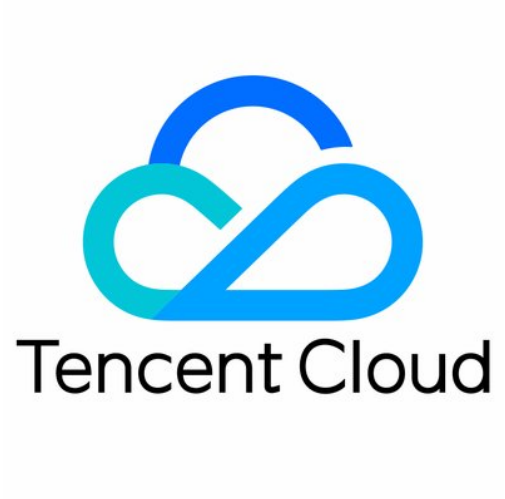

Contents

Migration to Cloud
- Feeling Migration
  - Metadata Migration
  - Message Service Data Flow Migration
- Seamless migration

Migration to Cloud
Feeling Migration
Metadata Migration

Last updated：2023-05-16 11:01:37

Overview

When using a TDMQ for RocketMQ cluster, you may need to migrate your existing business, for example, from a self-built or third-party RocketMQ cluster to the TDMQ for RocketMQ cluster.
This document describes how to migrate metadata from self-built open-source RocketMQ to TDMQ for RocketMQ.
Note
Both dedicated clusters and virtual clusters support metadata migration, and the migration steps are the same.

Prerequisites

You have created a cluster advance and a namespace in the console.

Directions

Step 1. Export the metadata file

If you are using self-built open-source RocketMQ, you can export metadata in the following two ways:

Option 1

If your RocketMQ server can access the public network, run the following script on your server directly (if there are multiple servers in your RocketMQ cluster, you can run the script on any server as long as the network is interconnected in the cluster).

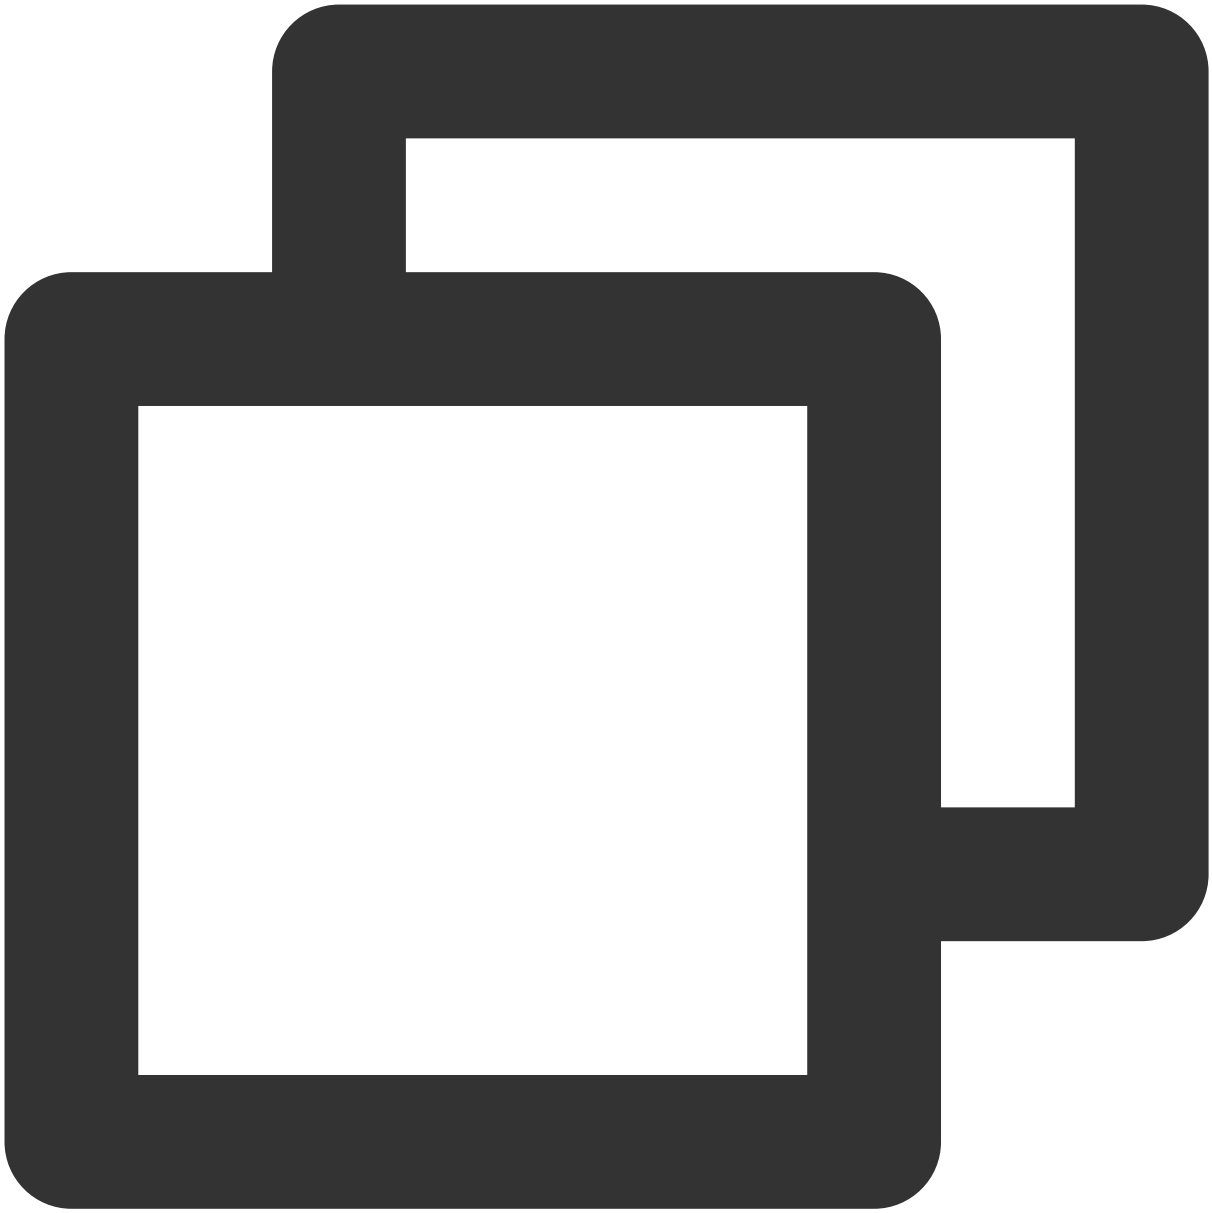

```
/bin/bash -c "$(curl -fsSL https://rocketmq-1306598660.cos.ap-guangzhou.myqcloud.co
```

Option 2

If your RocketMQ server has no permission to access the public network, follow the steps below:

1. Download the migration tool.

2. Upload the tool to your self-built RocketMQ cluster (if there are multiple servers in your RocketMQ cluster, you can run the tool on any server as long as the network is interconnected in the cluster).

3. Decompress the tool and enter the directory.

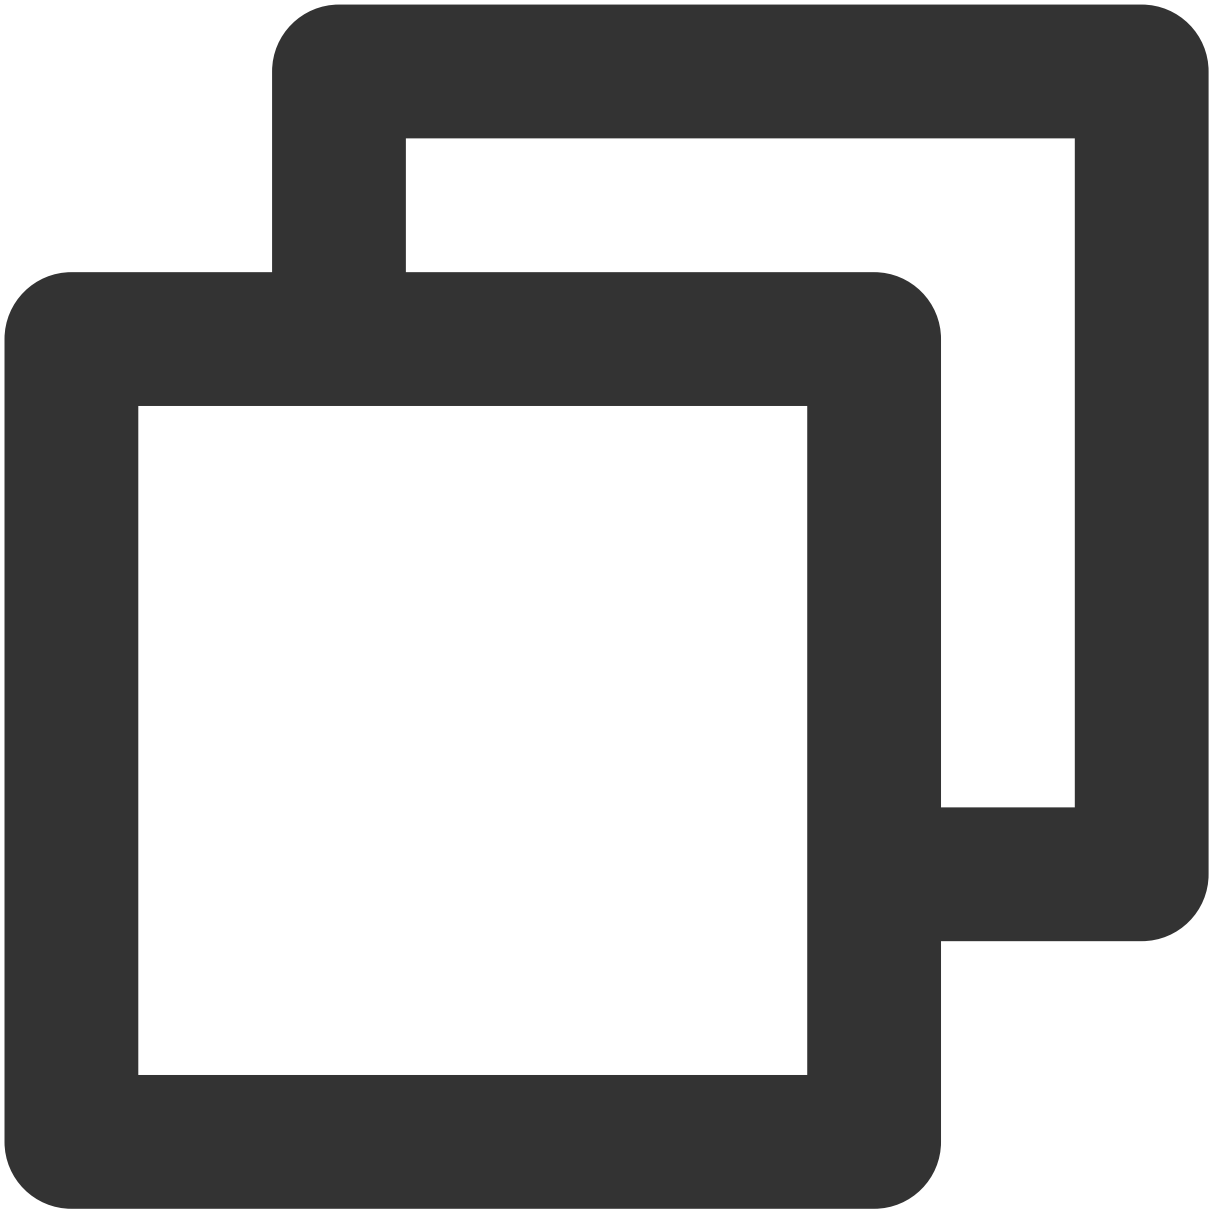

```
unzip rocketmq-migration.zip
cd rocketmq-migration
```

4. Run the following command for migration.

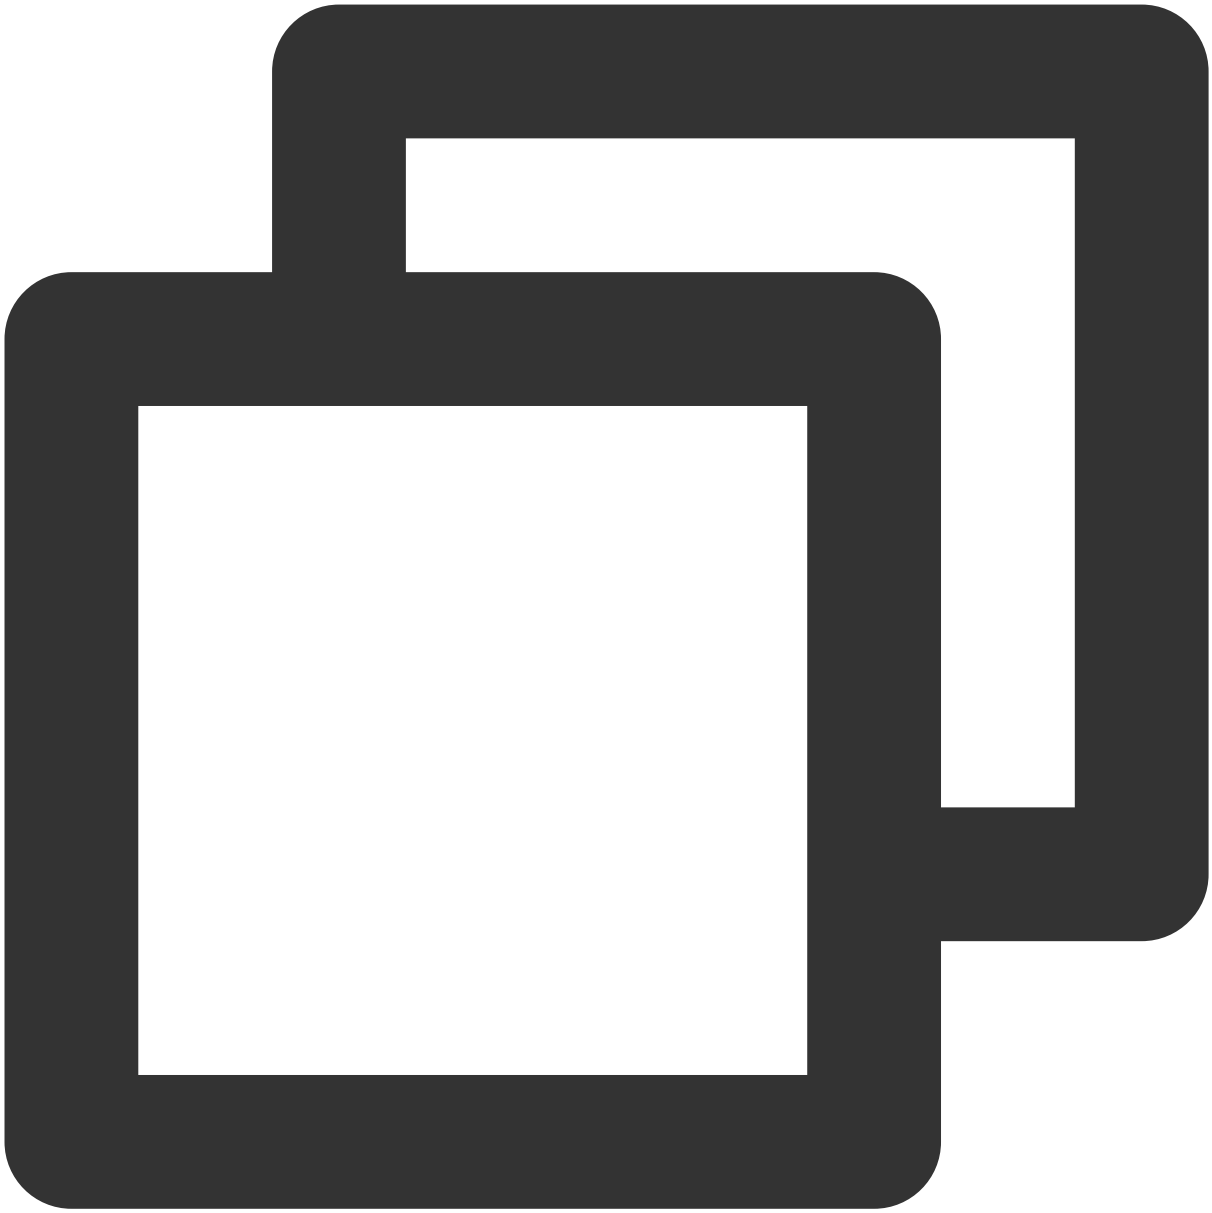

```
./bin/export.sh
// Enter the open-source RocketMQ address, such as `localhost:9876`
Enter name server address list:localhost:9876
// Select a cluster to export, such as `DefaultCluster`
Choose a cluster to export:DefaultCluster
// Enter a directory for saving the exported metadata, which is `/tmp/rocketmq/conf
Enter file path to export [default /tmp/rocketmq/export]:
```

Step 2. Create a migration task

1. Log in to the TDMQ for RocketMQ console and enter the **Migration to Cloud** page, enter the migration task list page, and click **Create Task** to create a corresponding migration task.
2. Select the migration task type:

Cluster migration: This type of task migrates metadata from the self-built RocketMQ cluster to the TDMQ for RocketMQ cluster. The migration tool will parse the part before "%" of each topic name in the open-source dashboard as the namespace name by default, so that you can create multiple logically isolated namespaces. If no topic name can be parsed in the self-built cluster, a namespace named `default` will be generated automatically.

Specified namespace import: This type of task migrates metadata from the self-built RocketMQ cluster to a specified TDMQ for RocketMQ namespace. If there are no namespaces in topics in the self-built cluster, you can select specific topics and groups you want to import, and specify the TDMQ for RocketMQ namespaces to which they are imported to distinguish between businesses or environments.

3. Upload the metadata file obtained in step 1 and select the topics and groups you want to import.

Note

Up to 1,000 topics and 1,000 groups can be imported in a single task. Excess data will fail to be imported.

Step 3. Check the task status

After the task is successfully created, enter the task list to view the task status. If there is too much data, the task needs to load for a while. Click **View Details** to view the specific running status of the task.

If the task status is **Some failed** or **All failed**, you can filter causes of the failures in the **Task Status** column.

Message Service Data Flow Migration

Last updated：2024-01-18 10:07:12

Overview

In a metadata migration task, you can sync the metadata of a self-built RocketMQ cluster to TDMQ for RocketMQ. After the sync, you need to modify the access information of the producer and consumer clusters in order to migrate them from the self-built cluster to TDMQ for RocketMQ for message sending and receiving.

Note:

Currently, the message migration service migrates only message production and consumption linkages but not message data in the old RocketMQ cluster. It applies only to exclusive clusters as the destination and will be supported for virtual clusters after the beta test ends.

Migration Directions

This document describes the **double-read double-write** and **batch release** schemes of the message migration service. During migration, the producer and consumer clusters can produce or consume messages in the old RocketMQ and new TDMQ for RocketMQ clusters in parallel. Data will not be heaped because of the migration, so the business can transition smoothly.

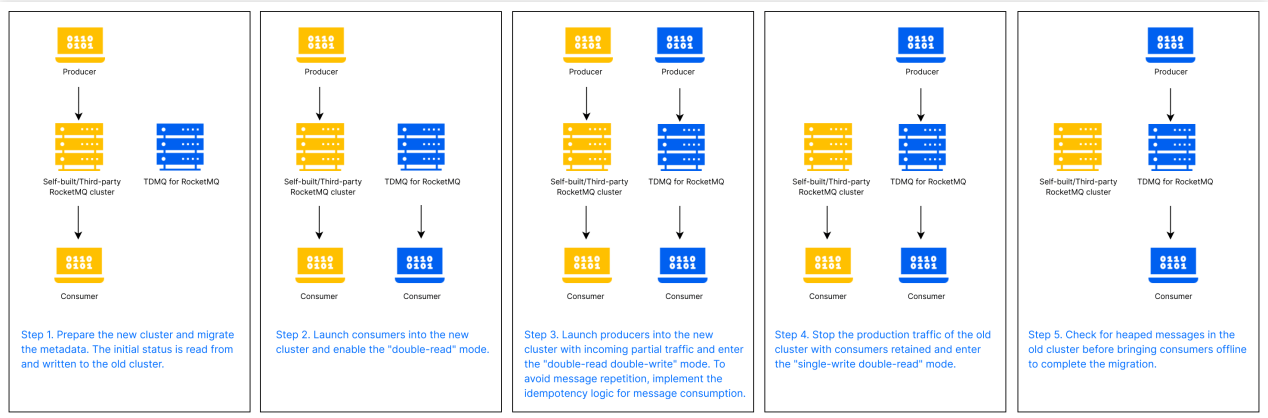

The detailed directions are as follows:

1. Create a TDMQ for RocketMQ cluster, migrate the metadata, and get the required client information in the console, such as the access point of the new cluster, `AccessKey` , and `SecretKey` .

2. Modify the access information of certain nodes in the consumer cluster to connect corresponding consumers to the new TDMQ for RocketMQ cluster. They will consume messages in the new cluster, while the rest will continue to

consume messages in the old cluster.
3. Modify the access information of certain nodes in the producer cluster to connect corresponding producers to the new TDMQ for RocketMQ cluster. They will send messages to the new cluster, while the rest will continue to send messages to the old cluster. To avoid message repetition or loss, implement the idempotency logic for message consumption in advance.
4. Connect the remaining producers to the new TDMQ for RocketMQ cluster. Then, all messages will be sent to the new cluster.
5. Check whether there are heaped messages that are not consumed in the old RocketMQ cluster, and if not, connect the remaining consumers to the new TDMQ for RocketMQ cluster. At this point, the migration is completed.

Note:

If you don't follow the above steps strictly, for example, if you switch producers first and then consumers, message loss may occur.
Before switching the remaining consumers, make sure that all messages in the old RocketMQ cluster have been consumed; otherwise, some messages may not be consumed. You can view the number of heaped messages in the old cluster to check whether consumption has completed.

Migration process diagram

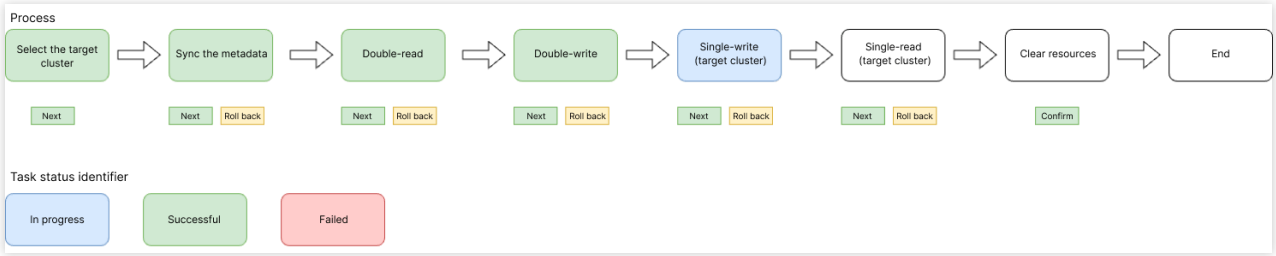

Possible Issues

Sequence

Message sequence cannot be guaranteed due to cluster switch.

Message repetition

Message repetition occurs only in extreme cases. For example, if a consumer consumes a message but doesn't send an acknowledgment to the server (the old RocketMQ cluster), the message will be put into the retry queue, causing repeated consumption. Implementing the idempotency logic can avoid this issue.

Consumption delay

During the read switch, partitions are reallocated with rebalancing between the queue and consumer client. This may cause a short consumption delay, but you don't need to handle it because it won’t persist after the switch.

Seamless migration

Last updated：2024-06-13 16:05:13

When using the TDMQ for RocketMQ, clients often need to switch their existing businesses. For example, from a self-built RocketMQ or another platform's managed RocketMQ to the TDMQ for RocketMQ.
In the process of migration, clients generally have two methods: noticeable migration and seamless migration.
Noticeable migration means that clients first import and export metadata, and then, following the Data Flow Migration Guide, switch the consumer and producer access addresses in sequence. However, this type of migration is somewhat invasive to clients' online services. During the switch of access points, it may affect the online business, thus it is more suitable for migration scenarios with low or no continuity requirements.

Overview

Currently, Tencent Cloud TDMQ for RocketMQ supports clients in migrating the entire cluster without interrupting business. It provides visual insights into the progress of migration tasks at different stages through the console. And it supports pausing and rolling back tasks.
This task guides you through the seamless migration of an open-source RocketMQ cluster to Tencent Cloud TDMQ for RocketMQ.

Migration Directions

Prerequisites

Before starting the migration, you need to select the task of type **Create Cluster for Seamless Migration** on the **Migration to Cloud** page, and fill in the task information as guided below. **Before creating the task, it is recommended to review the current traffic situation of the source cluster and purchase a suitable TDMQ for RocketMQ cluster** to avoid throttling due to the cloud cluster's specification being too small after migration.
Task name: Name the migration task according to the actual business situation.
Source cluster name: Name your self-built cluster according to the actual business situation.
Target cluster type: Currently supports RocketMQ 4.x exclusive cluster and RocketMQ 5.x cluster.
Target cluster region: Select according to actual business needs. Only regions with the target cluster available are displayed.
Target cluster: You can choose between RocketMQ 4.x generic clusters and exclusive clusters or RocketMQ 5.x clusters based on actual needs, ensuring that the target cluster's specifications can fully accommodate the source cluster's traffic.
Note:

Since the clusters adapted for migration to the cloud require a separate installation of migration components, it's necessary to separately purchase the 4.x exclusive clusters that support seamless migration. You can also enable the **Migrating to the Cloud Component** switch when buying the exclusive clusters, as shown below. RocketMQ 4.x generic clusters and the full series of 5.x are not affected by the migration components and the corresponding clusters you previously purchased can be directly used.

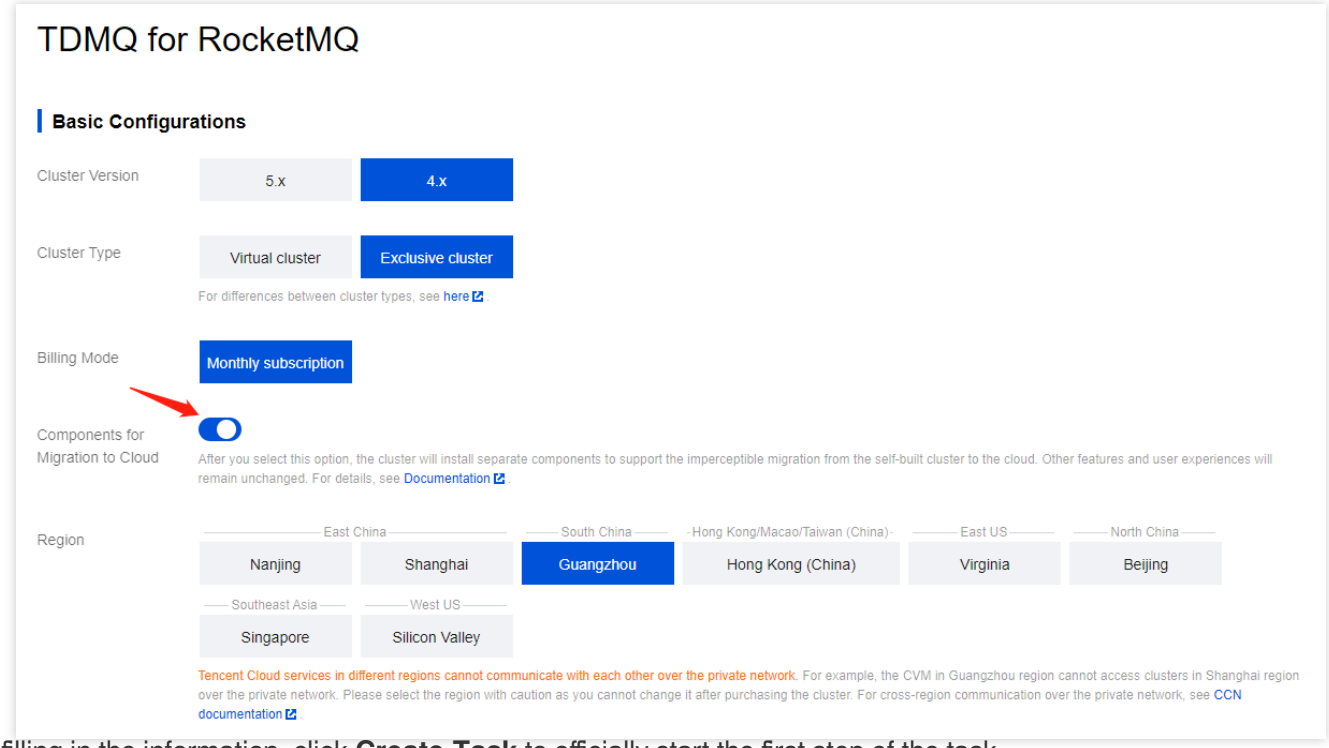

After filling in the information, click **Create Task** to officially start the first step of the task.

On the task details page, you can always view the progress of the tasks you have created and view the current architecture diagram.

Step 1: Connecting to the Source Cluster

In this step, you must confirm the type of network connection and fill in the relevant information for the source cluster: Network connection type: The type of network connection used when connecting a RocketMQ cluster on Tencent Cloud to a source cluster. Depending on the actual situation, you can choose between the public network and VPC network. Subsequent steps will verify the network connection. If you choose VPC, you will need to select a specific VPC and subnet.

Note:

If you choose to use the public network, to ensure the migration goes smoothly, you need to allow Tencent Cloud to access resources within your corresponding network. You can select your region in the table below and add Tencent Cloud's IP addresses for the corresponding regions to the public network allowlist of your self-built cluster.

Source Cluster's Located Region	Public IP Ranges to Be Opened
Beijing	152.136.0.0/16 81.70.0.0/16

Shanghai	121.4.0.0/16 106.54.0.0/16 49.234.0.0/16 49.235.0.0/16 175.24.0.0/16
Guangzhou	119.29.0.0/16 106.52.0.0/16 106.53.0.0/16 106.55.0.0/16 118.89.0.0/16 81.71.0.0/16 42.194.0.0/16
Nanjing	119.45.0.0/16 129.211.0.0/16 146.56.0.0/16 175.25.0.0/16
Hong Kong (China)	119.28.0.0/16 101.32.0.0/16
Singapore	43.134.0.0/16 101.32.0.0/16
Silicon Valley	170.106.0.0/16 49.51.0.0/16
Virginia	43.130.0.0/16 170.103.0.0/16
Shanghai Finance	212.129.0.0/16

NameServer address: Based on the network connection type entered in the previous step, fill in the NameServer address of the source cluster. The format is IP + port. Separate multiple addresses with semicolons.
NameServer address type: Only required when using VPC network to access the source cluster. Depending on different customers' habits, the NameServer address could be a CVM address or a CLB address on Tencent Cloud. Please fill in truthfully. This is only for identifying the network connectivity scenario and has no impact on subsequent operations.
Whether ACL enabled on the source cluster: Fill in truthfully. If your source cluster has ACL management for the Admin API enabled, to ensure smooth migration, you need to fill in the source cluster's accessKey and secretKey

(here it refers to the ak and sk for calling the Admin API, not the ak sk used for sending and receiving messages, which needs to be distinguished).

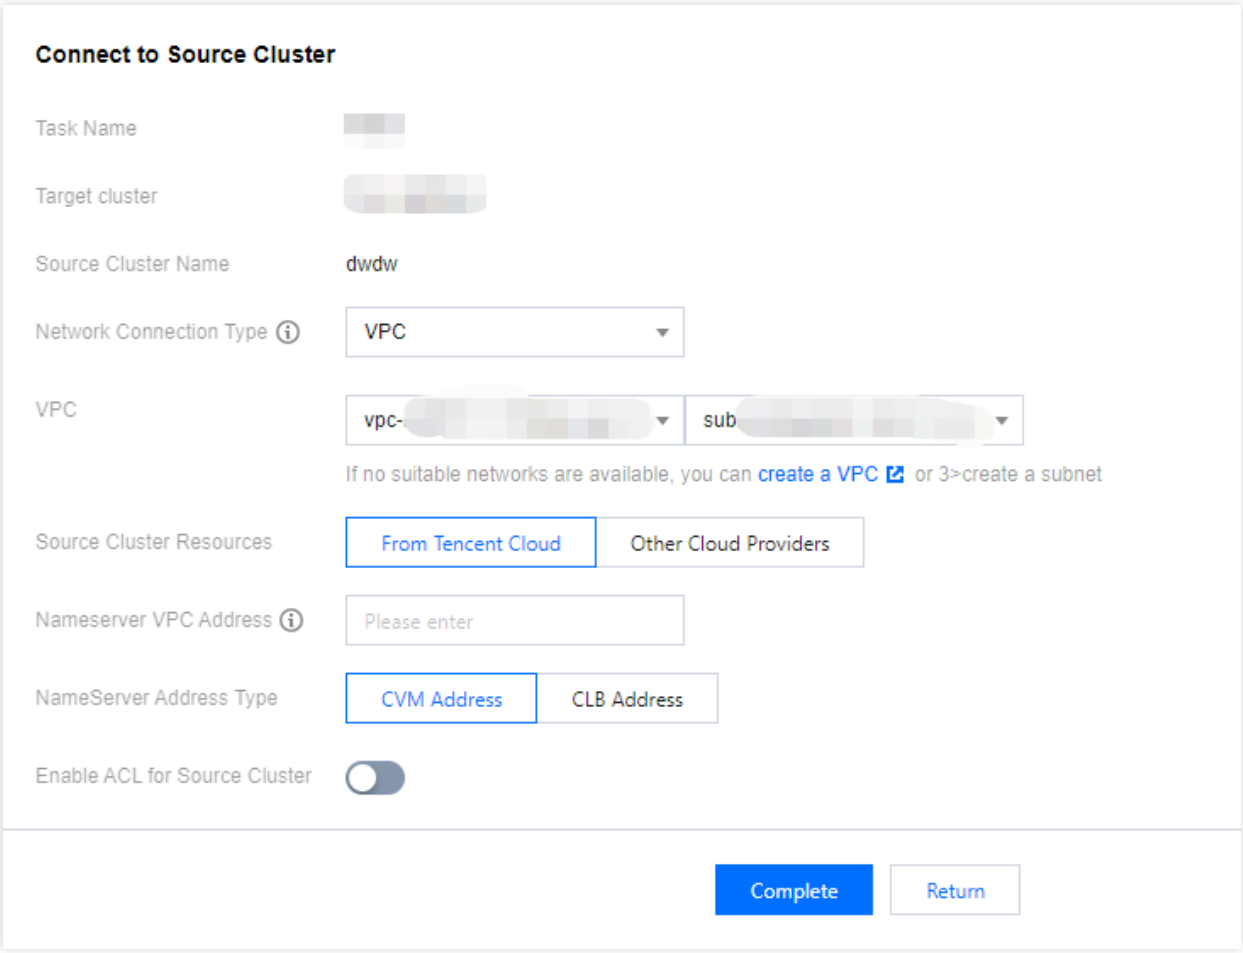

Note:

Depending on different customers' migration scenarios, there are situations where multiple source clusters migrate to the same TDMQ for RocketMQ cluster (target cluster). In such scenarios, the following issues need attention:

1. A target cluster can only be involved in one migration task at a time. Therefore, if the above situation occurs, please migrate multiple source clusters in batches.
2. When there are multiple migration tasks, different source clusters (suppose they are source cluster A and B) may have different ACL enabling configurations. Suppose during the first migration task, the source cluster A is migrated, and during the task creation stage, the ACL switch for source cluster A is enabled, then the ACL switch for the actual target cluster will also be enabled. During the second migration task, i.e., migrating source cluster B, the page will default to requiring ACL information when connecting to source cluster B. If the actual source cluster B does not use ACL, then in the subsequent traffic cut-over process, it will be necessary to supplement all clients of source cluster B with the relevant AK/SK configurations.
3. If in the above migration scenario, the source cluster A does not activate the ACL switch during the first migration task, then by default, the client of source cluster B will not be subject to AK/SK verification.

After filling in the above information, you can click the **Complete Network Configuration** at the bottom of the page. The TDMQ for RocketMQ migration component will connect to the source cluster. If the connection is successful, you

can proceed to the next step. If the connection failed, the page will display the reason for the error. Common reasons and solutions for such errors include:

Reason	Solution
Connecting to NameServer failed	Confirm if the NameServer address entered is correct. If you are using a public network connection, ensure the public network allowlist of the NameServer has added Tencent Cloud's IP to the allowlist, or check if the configuration of the allowlist is correct. If you are using a public network connection, ensure the target cluster, namely the Tencent Cloud RocketMQ cluster you purchased, has public network connections enabled.
Authentication failed	Confirm if the source cluster has enabled permission control. If yes, ensure the accessSecret and accessKey entered are correct.

Step 2: Importing Metadata

After successfully connecting to the NameServer of the source cluster, you can start syncing the metadata from the source cluster to the target cluster.

The migration tool automatically scans for related metadata of the source cluster, such as topics and groups. As shown below, you can confirm the information about topics and groups on the page. After the information is confirmed, you can click the **Operation** column's **Confirm and Import** button to proceed. You can also perform a batch import by clicking the

button in front of the list.

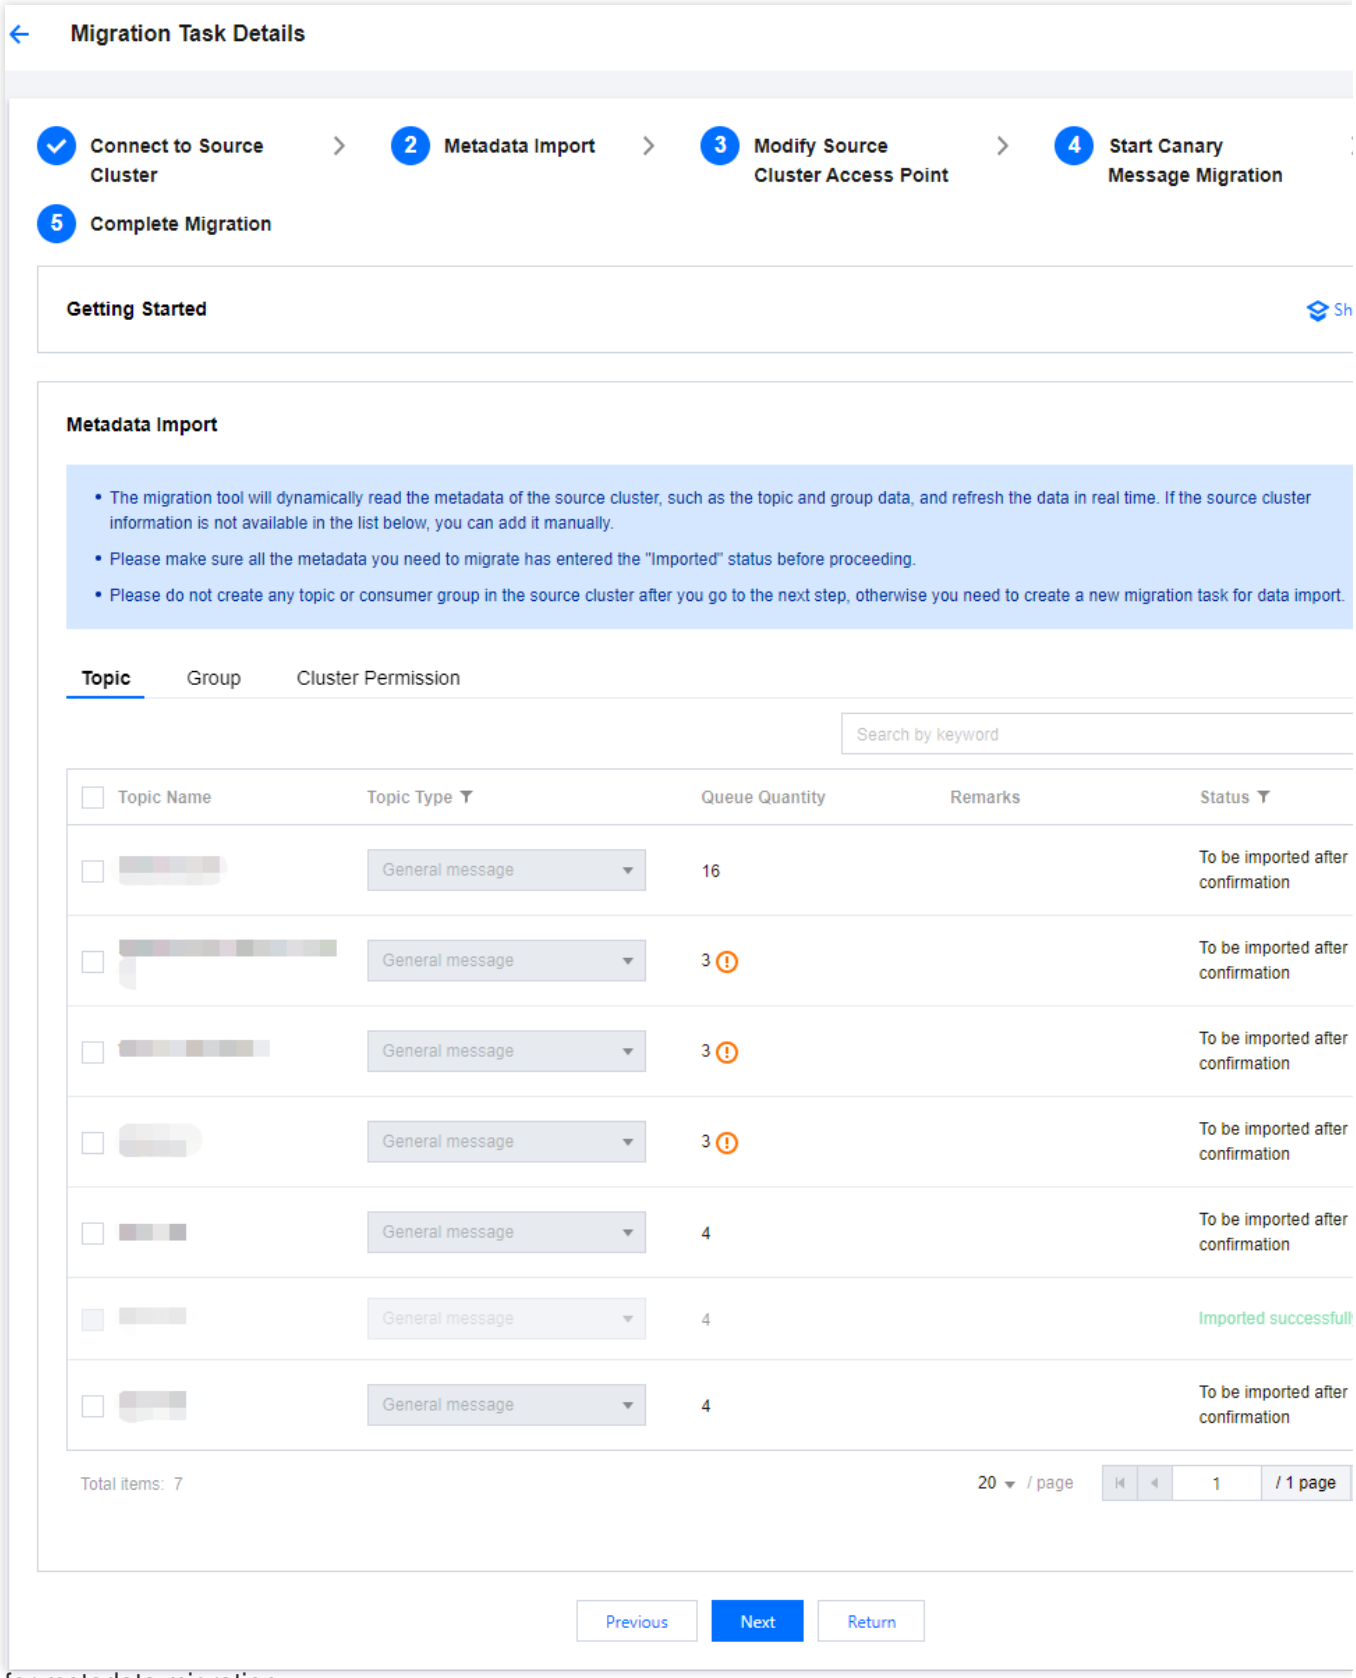

Notes for metadata migration:

If you cannot find the topics or groups created on the source cluster in the current list, you can click the **Refresh** button in the upper right corner to refresh.

If some topics and groups from the source cluster are still missing after refreshing, you can click the **Create** button to create new topics or groups.
The number of queues for a topic defaults to the same as that in the source cluster. For stability reasons, if the number of queues for a topic in your source cluster does not meet the TDMQ for RocketMQ range (for example, exceeding 16 queues or fewer than 3 queues), TDMQ for RocketMQ will fine-tune the number of queues, and a prompt will appear on the page.
The migrated groups default to the TCP protocol. If you need to create groups that support the HTTP protocol, you can do so separately on the target cluster page.
In situations where the target cluster is an exclusive 4.x cluster, if namespaces are used in the source cluster, the migration tool will, by default, identify and use the part of the topic or group name before the % character as the namespace name. If namespaces are not enabled in the source cluster, the migration tool will, by default, place all topics and groups into a predefined namespace named Default.
In the metadata import stage, you can also add the roles used for sending and receiving messages within the cluster, as well as the AK/SK. If there are many roles, you can click the

icon in the upper right corner for batch import.
After importing metadata, if you need to create new topics/groups, or operate on topics/groups, go to the target cluster, that is, the corresponding cluster's console in RocketMQ's console, to create.
Note:
If you encounter metadata format compatibility issues during migration, or the target cluster does not fully access the metadata information of your source cluster, you can contact us through a ticket.

Step 3: Modifying the Source Cluster's Endpoint

After it is confirmed that all the metadata to be migrated has been imported, you need to follow the instructions on the page and modify the endpoint addresses of all clients.
After modifying the endpoint addresses of clients, the migration tool will forward traffic to the source cluster, so the actual connections of each client remain with the source cluster, thus this step carries no risk. Before clicking Next, ensure that the endpoint of all clients (including consumers and producers) have been completely switched. Otherwise, during the grayscale migration of topics in the subsequent step, clients that have not changed their endpoint will be unable to connect to the target cluster, thus preventing them from participating in the grayscale migration.
You can identify whether all clients have completed the modification of the endpoint by viewing the number of clients connected to each topic and the most recent connection time displayed on the **Endpoint Modification Details** list page.
If all client endpoints have been successfully modified, you can proceed by clicking **Next** to officially enter the traffic switching stage.

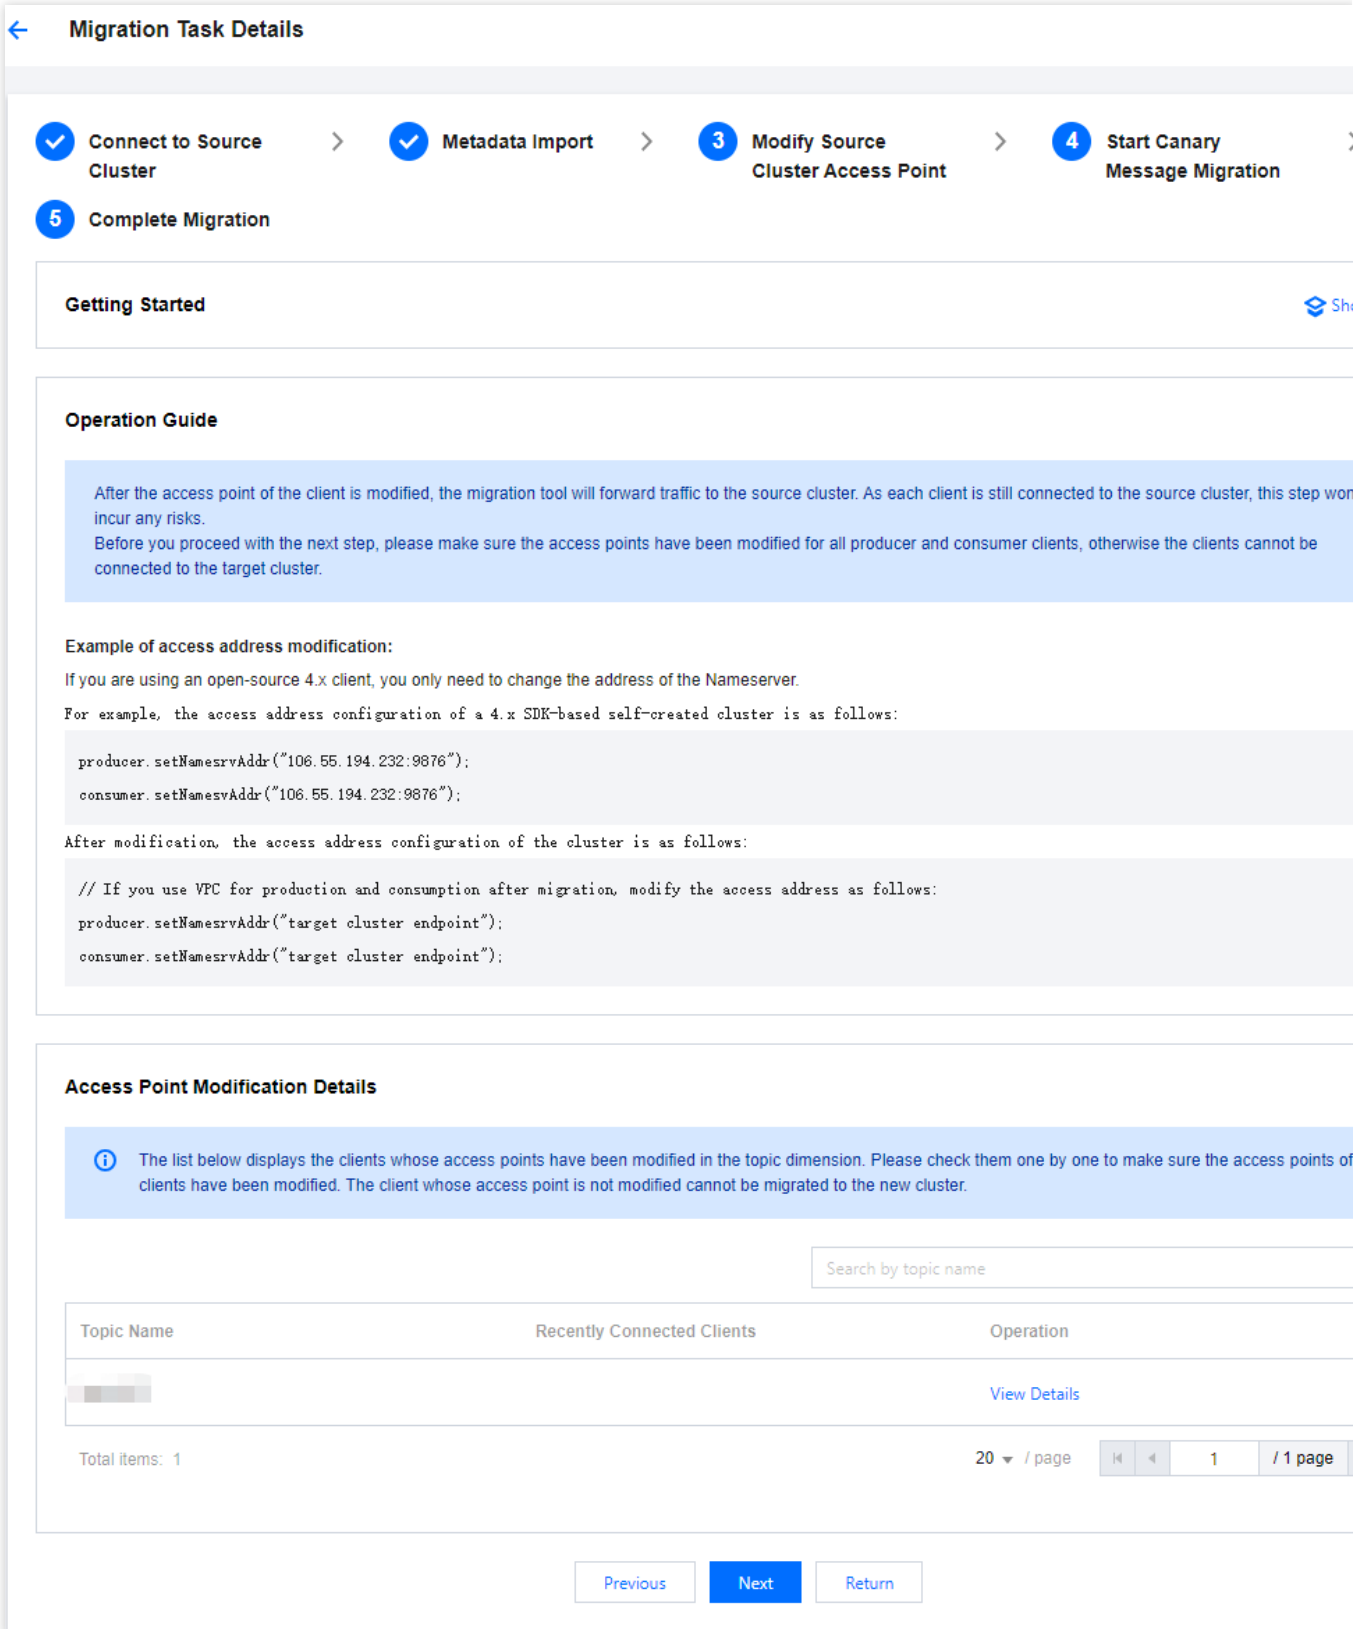

On the **Access Point Modification Details** list page, a comparison of the number of connected clients for each Group under the source and target clusters is displayed. This allows you to quickly locate any consumer client that has not completed modification. As shown below, the list page will display Groups that are inconsistent between the source and target clusters and automatically expand the Client lists on both sides for comparison. In the remark column, you can quickly locate the reasons.

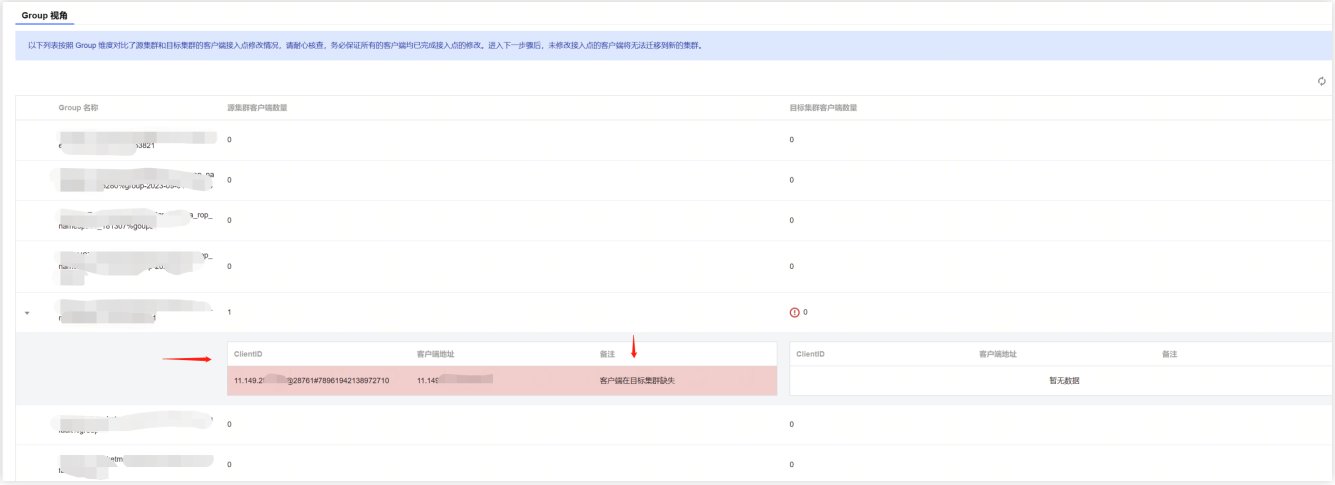

If all client access points have been confirmed to be modified, you can continue to click **Next** to officially enter the traffic switch phase. Before you enter the traffic switch phase, if any client access points are not yet modified, the page will give a prompt. If no popup appears, it can be assumed that all consumer client access points have been switched.

Step 4: Grayscale Migration of Messages

In the grayscale migration stage, the migration tool will migrate the migrants one by one at the topic granularity, following the sequence: Initial State (read and write on source cluster) > Dual Read Enabled (write to source cluster, read from both) > Dual Read and Write (read from both, write to both) > Cutting Over (write to target cluster, read from both) > Cut-over Completed (read and write on target cluster)". Throughout the migration process, each state allows for a rollback to the previous state:

Initial state: The state of reading and writing on the source cluster. It is the starting state of migration. The read and write traffic is proxied through the migration component and still accesses the source cluster, thus causing no intrusion to the business side.

Dual read enabled: The message producer client writes to the source cluster, while the message consumer simultaneously reads traffic from both the source and target clusters.

Dual read and write: Messages sent by the message producer client are routed to either the source cluster or the target cluster at random. You can view the traffic of different clusters on the monitoring page; meanwhile, the message consumer reads traffic from both the source and target clusters simultaneously.

Cutting over: The message producer client writes to the target cluster, while the message consumer reads traffic from both the source and target cluster simultaneously. You need to verify the new message transmission link is functioning normally without any anomalies and wait for the remaining messages in the source cluster to be fully consumed.

Cut-over Completed: After it is confirmed that the new message transmission link meets expectations in the previous stage, and under the circumstance that all messages have been consumed in the source cluster without any backlog, the system enters the state of reading and writing on the target cluster, with all traffic accessing only the new target cluster.

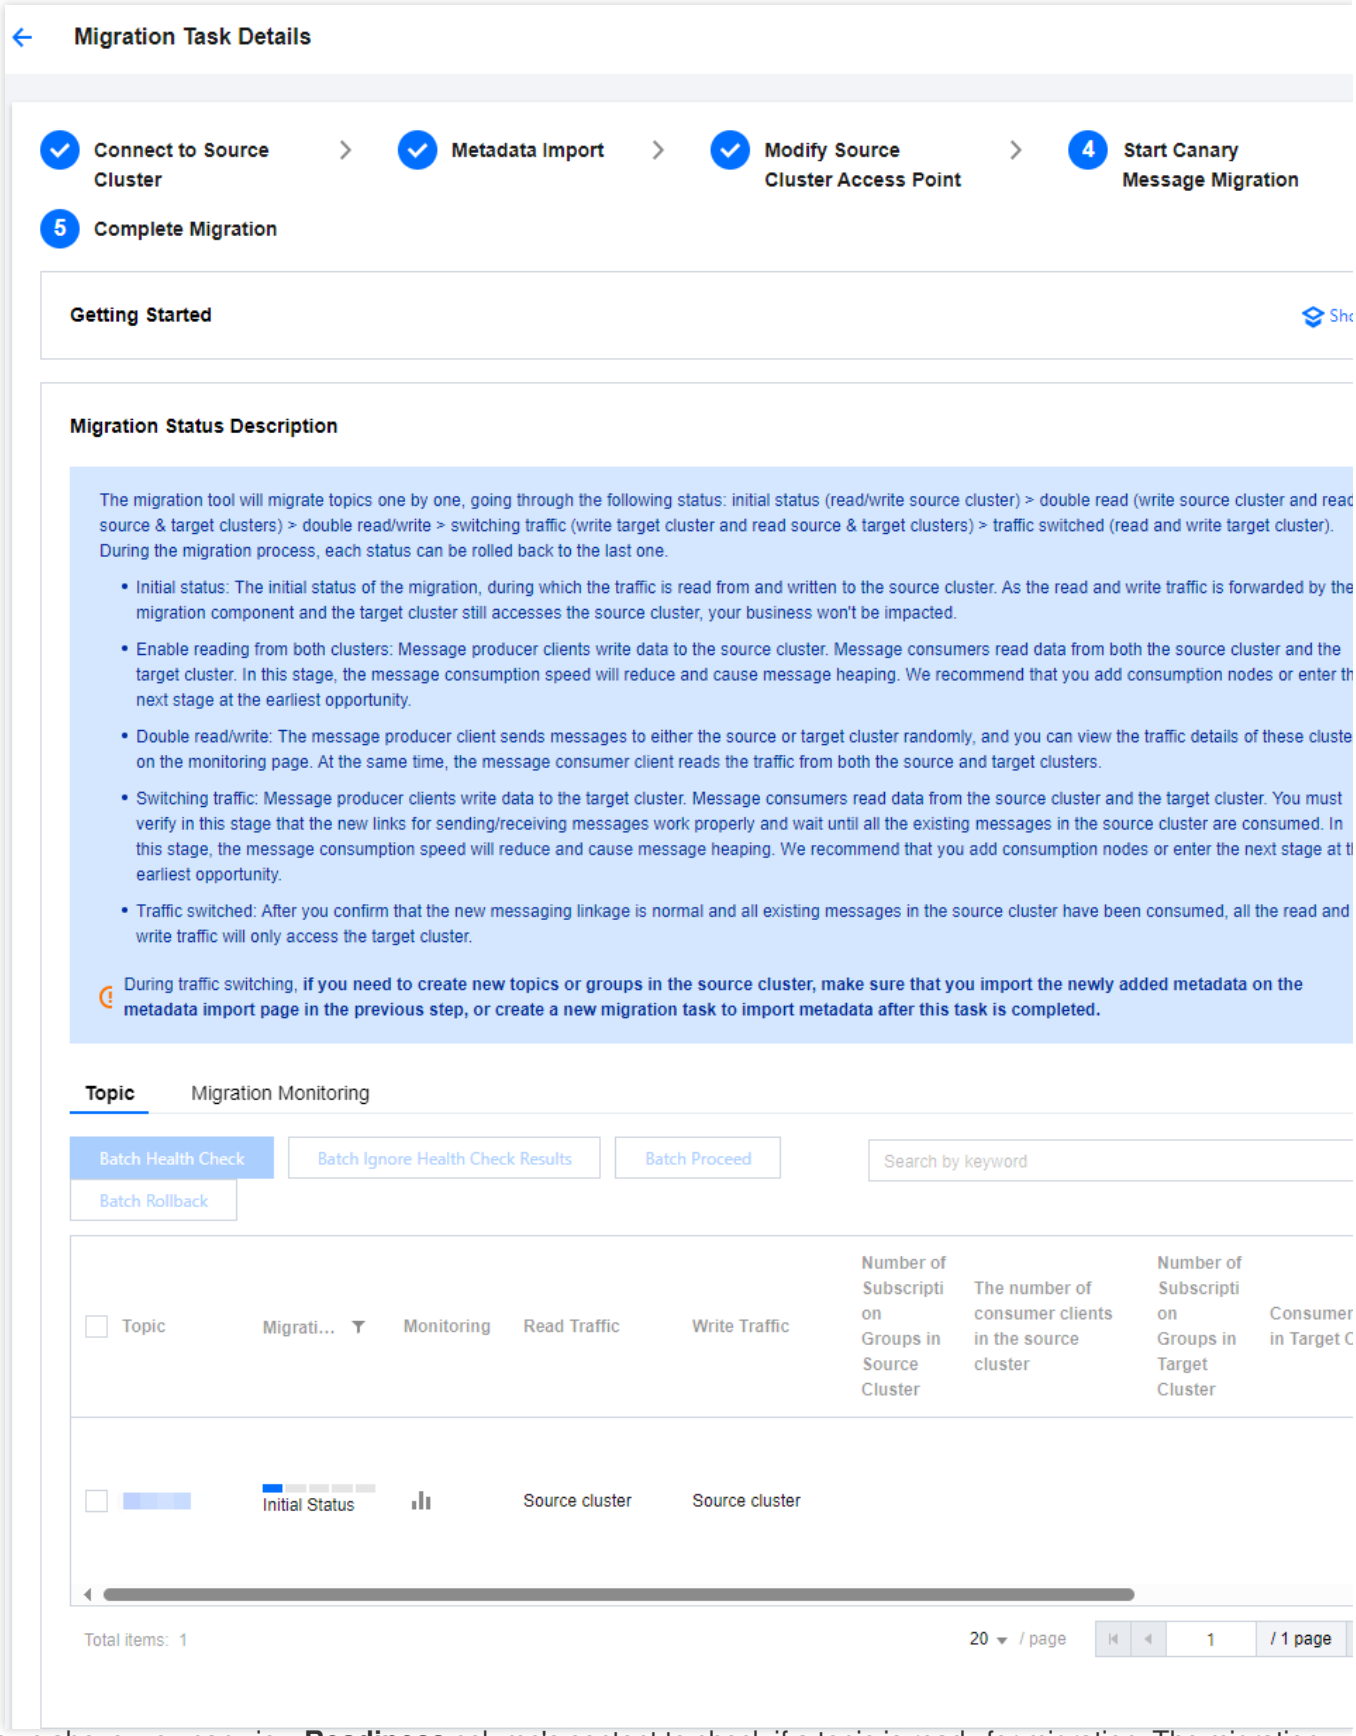

As shown above, you can view **Readiness** column's content to check if a topic is ready for migration. The migration tool will perform batch scans periodically. You can also use the **Operation** column's **Health Check** button for real-

time checks on an individual topic. Topics that is in the Ready state can proceed to the next stage. During the cut-over process, you can click the **Rollback** to rollback an individual topic or batch.

During the topic cut-over process, you can view the entire cluster or an individual topic's traffic and operational state through the **Migration Monitoring** Tab. Only topics in the Initial State support being removed from the migration task.

During the cut-over process, you can freely compare the target cluster and source cluster's Group (consumer group) and the number of clients, as shown above, and it will also indicate which Groups in the source cluster have not yet completed the switch.

During the cut-over process, you can also view the traffic changes between the source and target clusters via the monitoring tab. The monitoring data will be displayed by using different legends.

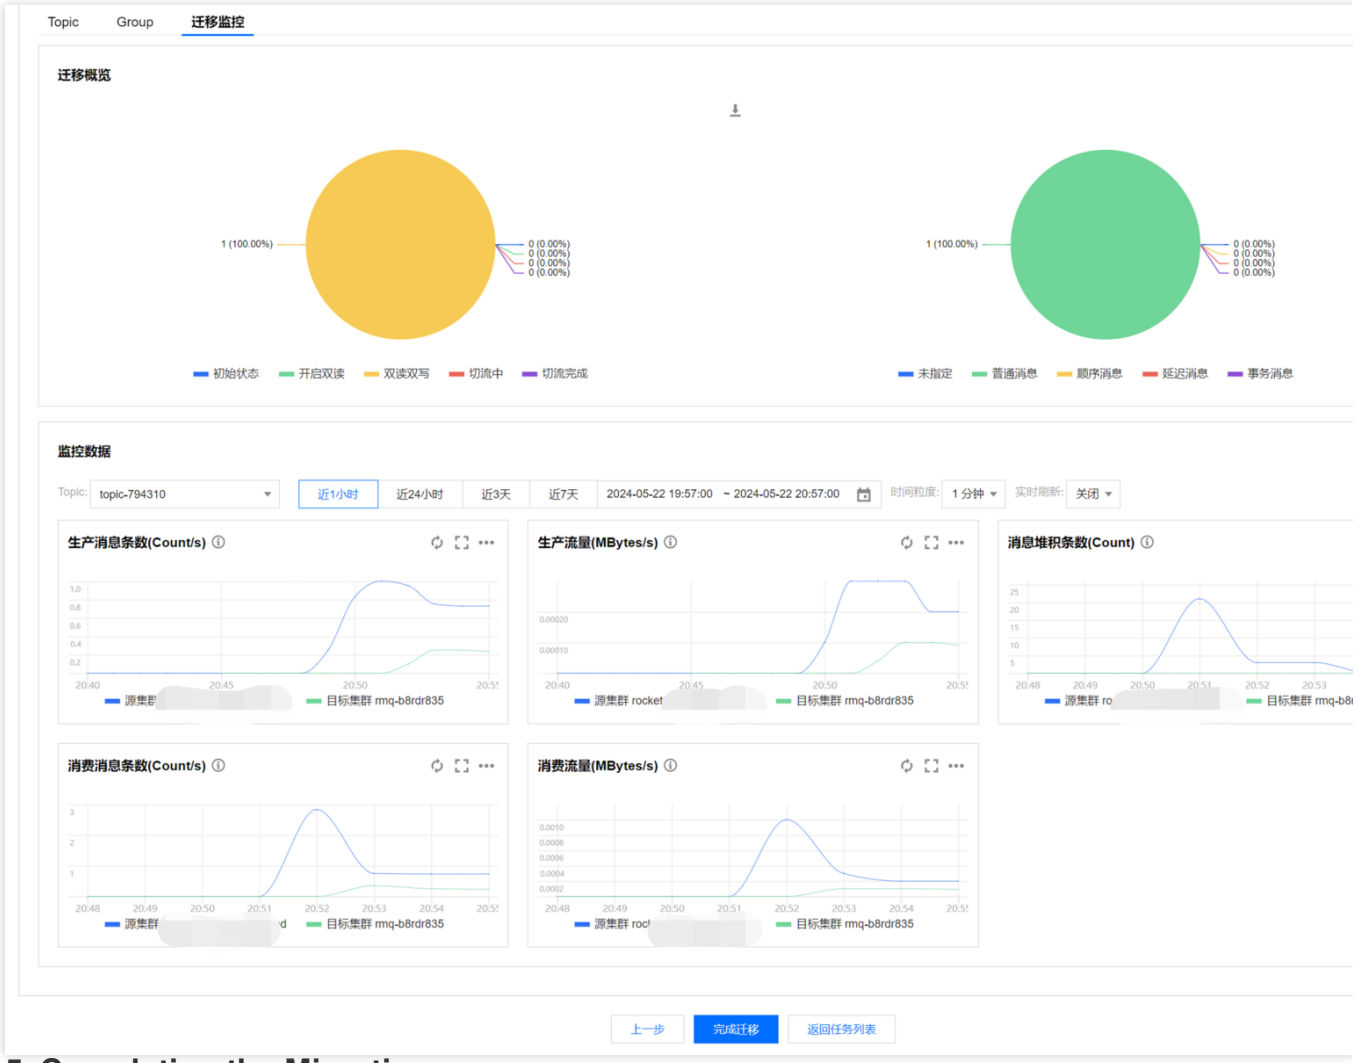

Step 5: Completing the Migration

Once all topics are in the Cut-over Completed state, you can proceed to the next step, which completes the entire migration process. After the migration is completed, all topics and clients are migrated, and all message traffic runs on TDMQ for RocketMQ. You can monitor the target and source clusters to gradually decommission the source cluster.