

Data Lake Insight

User Guide

Issue 01
Date 2023-12-11



Copyright © Huawei Technologies Co., Ltd. 2023. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to "Vul. Response Process". For details about the policy, see the following website:<https://www.huawei.com/en/psirt/vul-response-process>
For enterprise customers who need to obtain vulnerability information, visit:<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 DLI Console Overview.....	1
2 SQL Editor.....	2
3 Job Management.....	9
3.1 Overview.....	9
3.2 SQL Job Management.....	10
3.3 Flink Job Management.....	15
3.3.1 Overview.....	15
3.3.2 Managing Flink Job Permissions.....	18
3.3.3 Preparing Flink Job Data.....	21
3.3.4 (Recommended) Creating a Flink OpenSource SQL Job.....	22
3.3.5 Creating a Flink SQL Job.....	31
3.3.6 Creating a Flink Jar Job.....	37
3.3.7 Performing Operations on a Flink Job.....	43
3.3.8 Flink Job Details.....	50
3.3.9 Tag Management.....	57
3.4 Spark Job Management.....	60
3.4.1 Spark Job Management.....	60
3.4.2 Creating a Spark Job.....	61
4 Queue Management.....	69
4.1 Overview.....	69
4.2 Queue Permission Management.....	75
4.3 Creating a Queue.....	77
4.4 Deleting a Queue.....	80
4.5 Allocating a Queue to an Enterprise Project.....	80
4.6 Modifying the CIDR Block.....	81
4.7 Elastic Queue Scaling.....	82
4.8 Scheduling CU Changes.....	84
4.9 Testing Address Connectivity.....	88
4.10 Creating an SMN Topic.....	88
4.11 Managing Queue Tags.....	90
4.12 Setting Queue Properties.....	92
5 Data Management.....	94

5.1 Databases and Tables.....	94
5.1.1 Overview.....	94
5.1.2 Managing Database Permissions.....	97
5.1.3 Managing Table Permissions.....	103
5.1.4 Creating a Database or a Table.....	111
5.1.5 Deleting a Database or a Table.....	120
5.1.6 Modifying the Owners of Databases and Tables.....	121
5.1.7 Importing Data to the Table.....	121
5.1.8 Exporting Data from DLI to OBS.....	125
5.1.9 Viewing Metadata.....	128
5.1.10 Previewing Data.....	129
5.1.11 Managing Tags.....	129
5.2 Package Management.....	133
5.2.1 Overview.....	133
5.2.2 Managing Permissions on Packages and Package Groups.....	134
5.2.3 Creating a Package.....	138
5.2.4 Deleting a Package.....	140
5.2.5 Modifying the Owner.....	140
5.2.6 Built-in Dependencies.....	141
6 Job Templates.....	162
6.1 Managing SQL Templates.....	162
6.2 Managing Flink Templates.....	167
6.3 Managing Spark SQL Templates.....	172
6.4 Appendix.....	173
6.4.1 TPC-H Sample Data in the SQL Template.....	173
7 Enhanced Datasource Connections.....	176
7.1 Overview.....	176
7.2 Cross-Source Analysis Development Methods.....	178
7.3 Creating an Enhanced Datasource Connection.....	180
7.4 Deleting an Enhanced Datasource Connection.....	186
7.5 Modifying Host Information.....	186
7.6 Binding and Unbinding a Queue.....	188
7.7 Adding a Route.....	190
7.8 Deleting a Route.....	191
7.9 Enhanced Connection Permission Management.....	192
7.10 Enhanced Datasource Connection Tag Management.....	193
8 Datasource Authentication.....	196
8.1 Introduction.....	196
8.2 Creating a CSS Datasource Authentication.....	198
8.3 Creating a Kerberos Datasource Authentication.....	200
8.4 Creating a Kafka_SSL Datasource Authentication.....	204

8.5 Creating a Password Datasource Authentication.....	207
8.6 Datasource Authentication Permission Management.....	210
9 Global Configuration.....	212
9.1 Global Variables.....	212
9.2 Permission Management for Global Variables.....	213
9.3 Project Configurations.....	215
9.4 Service Authorization.....	216
10 Custom Images.....	219
10.1 Overview.....	219
10.2 Downloading a Basic Image.....	219
10.3 Using Custom Images.....	220
10.4 Basic Image Components.....	224
11 Permissions Management.....	226
11.1 Overview.....	226
11.2 Creating an IAM User and Granting Permissions.....	230
11.3 Creating a Custom Policy.....	232
11.4 DLI Resources.....	238
11.5 DLI Request Conditions.....	238
11.6 Common Operations Supported by DLI System Policy.....	239
12 Other Common Operations.....	244
12.1 Importing Data to a DLI Table.....	244
12.2 Viewing Monitoring Metrics.....	244
12.3 DLI Operations That Can Be Recorded by CTS.....	249
12.4 Quotas.....	253
A Change History.....	255

1 DLI Console Overview

The **Overview** page of the DLI console provides you with the DLI workflow and resource usage.

How to Use DLI

The process of using DLI is as follows:

1. **Create a queue.**

Queues are DLI's compute resources. There are SQL queues and general-purpose queues. For a SQL queue, you can only submit Spark SQL jobs. For a general-purpose queue, you can submit Spark programs, Flink OpenSource SQL jobs, and Flink Jar jobs.

2. **Prepare data.**

Create databases and tables before you run a Spark SQL job. Upload a program package before you run a Spark job or a Flink Jar job.

3. **Edit and submit a job.**

After you set the job parameters, you can submit the job.

4. **View job status.**

Go to the **Job Management** page to view the job execution status.

Queue Usage (hours)

The overview page provides usage durations of all queues or a specific queue.

- Usage Usage (hours): an overview of the usage of all resources.
- Selected Queue Usage (hours): recent usage of a single queue.

2 SQL Editor

Introduction

You can edit and run SQL statements in the SQL job editor to execute data query.

The editor supports SQL:2003 and is compatible with Spark SQL. For details about the syntax, see [Data Lake Insight SQL Syntax Reference](#).

To access the SQL editor, choose **SQL Editor** in the left navigation pane of the DLI console, or click **Create Job** in the upper right corner of the **Job Management > SQL Jobs** page.

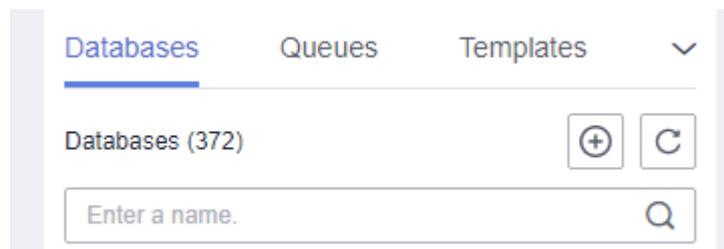
This topic describes the main functions of the SQL editor.

Notes

- If you access the SQL editor for the first time, the system prompts you to set a bucket for DLI jobs. The created bucket is used to store temporary data generated by DLI, such as job logs.
You cannot view job logs if you choose not to create the bucket. The bucket name will be set by the system.
On the OBS console, you can configure lifecycle rules for a bucket to periodically delete objects in it or change object storage classes.
- SQL statements can be executed in batches on the SQL editor page.
- Commonly used keywords in the job editing window are highlighted in different colors.
- Both single-line comment and multi-line comment are allowed. Use two consecutive hyphens (--) in each line to comment your statements.

Navigation pane

The navigation pane on the left consists of **Databases**, **Queues**, and **Templates** tabs.

Figure 2-1 SQL editor navigation pane**Table 2-1** Navigation pane buttons

No.	Name	Description
1	Databases	Displays all the existing databases and tables in these databases. <ul style="list-style-type: none">Click a database name to view the tables in the database.Click a table name to view the metadata in the table. A maximum of 20 metadata records can be displayed.After you double-click a table name, a SQL query statement is automatically entered in the editing window.
2	Queues	Displays existing queues.
3	Templates	Click the drop-down button to view 22 built-in standard TPC-H query templates and custom templates.

SQL Editing Window

SQL job editing window is displayed in the upper right part of the page.

The SQL statement editing area is below the operation bar. For details about keyboard shortcuts, see [Table 2-3](#).

Table 2-2 Components of the SQL job editing window

No.	Button & Drop-Down List	Description
1	Engine	SQL jobs support the Spark and Trino engines. <ul style="list-style-type: none">Spark is applicable to offline analysis.Trino is applicable to interactive analysis. Currently, only the default queue and Hive table query are supported.
2	Queues	Select a queue from the drop-down list box. If no queue is available, the default queue is displayed. Refer to Creating a Queue and create a queue. SQL jobs can be executed only on SQL queues.

No.	Button & Drop-Down List	Description
3	Database	Select a database from the drop-down list box. If no database is available, the default database is displayed. For details about how to create a database, see Creating a Database or a Table . NOTE If you specify the database in the SQL statements, the database you choose from the drop-down list will not be used.
4	Execute	Click this button to run the SQL statements in the job editing window.
5	Format	Click this button to format the SQL statements.
6	Syntax Reference	Click this button to view the <i>Data Lake Insight SQL Syntax Reference</i> .
7	Settings	Add parameters and tags. Parameter Settings: Set parameters in key/value format for SQL jobs. Tags: Set tags in key/value format for SQL jobs.
8	More	The drop-down list includes the following options: <ul style="list-style-type: none">Click Verify Syntax to check whether the SQL statements are correct.Click Set as Template to set SQL statements as a template. For details, see Managing SQL Templates.Click Change Theme to switch between dark and light modes.

Table 2-3 Keyboard shortcut description

Shortcut	Description
Ctrl+Enter	Execute SQL statements. You can run SQL statements by pressing Ctrl+R or Ctrl + Enter on the keyboard.
Ctrl+F	Search for SQL statements. You can press Ctrl+F to search for a required SQL statement.
Shift+Alt+F	Format SQL statements. You can press Shift + Alt + F to format a SQL statement.
Ctrl+Q	Syntax verification. You can press Ctrl + Q to verify the syntax of SQL statements.

Shortcut	Description
F11	Full screen. You can press F11 to display the SQL Job Editor window in full screen. Press F11 again to leave the full screen.

Executed Queries (Last Day) and View Result

After the SQL job is executed, you can view the execution history and result in the lower part of the editing area.

- **Executed Queries (Last Day)**

You can filter the execution history in the following ways:

- In the search box in the upper right corner of the **Executed Queries (Last Day)** pane, select a queue name or enter an execution statement in the search box.
- In the list, click the icon next to **Created** and choose **Ascending** or **Descending**.
- Select a job status from the **Status** list.

Table 2-4 Area description

Area	Description
<p>Executed Queries (Last Day)</p>	<p>The latest daily information about the submitted jobs, including the following items:</p> <ul style="list-style-type: none"> ● Queues: Queue name ● Username: User who executes the SQL statements ● Type: Type of the SQL job ● Status: Execution status of the SQL job ● Query ● Created ● Operation <ul style="list-style-type: none"> - Edit: Edit the SQL statement. - SparkUI: Switch to the SparkUI page to view the SQL statement execution process. <p>NOTE When you execute a job on a created queue, the cluster is restarted. It takes about 10 minutes. If you click SparkUI before the cluster is created, an empty projectID will be cached. The SparkUI page cannot be displayed. You are advised to use a dedicated queue so that the cluster will not be released. Alternatively, wait for a while after the job is submitted (the cluster is created), and then check SparkUI.</p> <p>Currently, only the latest 100 job information records are displayed on the SparkUI of DLI.</p> <p>This function is not supported for synchronization jobs and jobs running on the default queue.</p> <ul style="list-style-type: none"> - More: The following operations vary depending on the SQL job types and running status. <ul style="list-style-type: none"> Cancel: Cancel a SQL job that is running or being submitted. Re-execute: Execute the SQL statement again. View Result: View the execution result of a QUERY job. Export Result: Export the execution results of a QUERY job to a specified OBS path. View Log: View the OBS path for storing SQL statement execution logs. <p>NOTE The View Log button is not available for synchronization jobs and jobs running on the default queue.</p>

- **View Result**

Table 2-5 Operations in the result tab

Operation	Description
Clear the result	Clear the displayed SQL statement query results.
View chart/ table	Click  to view the query result in a chart or table.
Export the result	Click  to export the query result to OBS. For details, see Exporting Query Results . A maximum of 1000 records can be displayed in the query result on the console. To view more or all data, you can click Export Result to export the data to OBS.

SQL Query Procedure

1. Log in to the DLI management console. On the page displayed, choose **Job Management > SQL Jobs**. On the page displayed, click **Create Job**.

NOTE

On the SQL editor page, the system prompts you to create an OBS bucket to store temporary data generated by DLI jobs. In the **Set Job Bucket** dialog box, click **Setting**. On the page displayed, click the edit button in the upper right corner of the job bucket card. In the **Set Job Bucket** dialog box displayed, enter the job bucket path and click **OK**.

2. Select a queue from the queue list in the upper left corner of the SQL job editing window. For details about how to create a queue, see [Creating a Queue](#).
3. In the upper right corner of the SQL job editing window, select a database, for example, **qw**, from the **Databases** drop-down list.
4. Create a table, for example, **qw**. For details about how to create a database and table, see [Creating a Database or a Table](#).
5. In the SQL job editing window, enter the following SQL statement:

```
SELECT * FROM qw.qw LIMIT 10;
```

Alternatively, you can double-click the table name **qw**. The query statement is automatically entered in the SQL job editing window.
6. On top of the editing window, click **More > Verify Syntax** to check whether the SQL statement is correct.
 - a. If the verification fails, check the SQL statement syntax by referring to [Data Lake Insight SQL Syntax Reference](#).
 - b. If the syntax verification is successful, click **Execute**. Read and agree to the privacy agreement. Click **OK** to execute the SQL statement.
 - c. After the execution is complete, you can view the execution result in the area under the SQL job editing window.

7. (Optional) A maximum of 1000 records can be displayed in the query result on the current console. To view more or all data, click  to export the data to OBS.
8. (Optional) In the **View Result** tab, click  to display the query result in a chart. Click  to switch back to the table view.

 **NOTE**

- If no column of the numeric type is displayed in the execution result, the result cannot be represented in charts.
- You can view the data in a bar chart, line chart, or fan chart.
- In the bar chart and line chart, the X axis can be any column, while the Y axis can only be columns of the numeric type. The fan chart displays the corresponding legends and indicators.

Quickly Importing SQL Statements

- Double-click a table name in the navigation pane on the left to import the query statement of the selected table into the SQL statement editing window, and then click **Execute** to query.
- You can click **More** and choose **Save as Template** to save the SQL statement as a template for future use.

To use the SQL statement template, click **Templates** from the left pane of the SQL editor page. Double-click the required template in the template list, and modify it as required before executing the SQL statements.

3 Job Management

3.1 Overview

DLI Job Type

DLI provides the following job types:

- **SQL job:** SQL jobs provide you with standard SQL statements and are compatible with Spark SQL and Presto SQL (based on Presto). You can query and analyze heterogeneous data sources on the cloud through visualized APIs, JDBC, ODBC, or Beeline. SQL jobs are compatible with mainstream data formats such as CSV, JSON, Parquet, Carbon, and ORC.
- **Flink job:** Flink jobs are real-time streaming big data analysis service jobs running on the public cloud. In full hosting mode, you only need to focus on Stream SQL services and execute jobs instantly without being aware of compute clusters. Flink jobs are fully compatible with Apache Flink APIs.
- **Spark job:** Spark jobs provide fully-managed Spark compute services. You can submit jobs through the GUI or RESTful APIs. Full-stack Spark jobs, such as Spark Core, DataSet, Streaming, MLLib, and GraphX jobs, are supported.

Constraints

- DLI supports the following types of jobs: Spark SQL, Spark Jar, Flink SQL, and Flink Jar.
- DLI supports the following Spark versions: Spark 3.1, 2.4, and 2.3 (end of marketing, EOM).
- DLI supports the following Flink versions: Flink Jar 1.15, Flink 1.12 (EOM), Flink 1.10 (EOM), and Flink 1.7 (end of service or support, EOS).
- SQL jobs support the Spark and Trino engines.
 - **Spark:** displays jobs whose execution engine is Spark.
 - **Trino:** displays jobs whose execution engine is Trino.
- Only the latest 100 jobs are displayed on DLI's SparkUI.
- A maximum of 1,000 job results can be displayed on the console. To view more or all jobs, export the job data to OBS.

- To export job run logs, you must have the permission to access OBS buckets. You need to configure a DLI job bucket on the **Global Configuration > Project** page in advance.
- The **View Log** button is not available for synchronization jobs and jobs running on the default queue.
- Only Spark jobs support custom images.
- An elastic resource pool supports a maximum of 32,000 CUs.
- Minimum CUs of a queue that can be created in an elastic resource pool:
 - General purpose queue: 4 CUs
 - SQL queue: Spark SQL queue: 8 CUs; Trino SQL queue: 16 CUs

3.2 SQL Job Management

SQL jobs allow you to execute SQL statements entered in the [SQL job editing window](#), import data, and export data.

SQL job management provides the following functions:

- **Searching for Jobs:** Search for jobs that meet the search criteria.
- **Viewing Job Details:** Display job details.
- **Terminating a Job:** Stop a job in the **Submitting** or **Running** status.
- **Exporting Query Results:** A maximum of 1000 records can be displayed in the query result on the console. To view more or all data, you can export the data to OBS.

SQL Jobs Page

On the **Overview** page of the DLI console, click **SQL Jobs** to go to the SQL job management page. Alternatively, you can click **Job Management > SQL Jobs**. The job list displays all SQL jobs. If there are a large number of jobs, they will be displayed on multiple pages. You can switch to the specified page as needed. DLI allows you to view jobs in all statuses. By default, jobs in the job list are displayed in descending order of the job creation time.

Table 3-1 SQL Job management parameters

Parameter	Description
Queues	Name of the queue to which a job belongs
Engine	SQL jobs support the Spark and Trino engines. <ul style="list-style-type: none">• spark: displays jobs whose execution engine is spark.• trino: displays jobs whose execution engine is trino.
Username	Name of the user who executed the job.

Parameter	Description
Type	Job type. The following types are supported: <ul style="list-style-type: none">● IMPORT: A job that imports data to DLI● EXPORT: A job that exports data from DLI● DCL: Conventional DCLs and operations related to queue permissions● DDL: Conventional DDLs, including creating and deleting databases and tables● QUERY: A job that queries data by running SQL statements● INSERT: A job that inserts data by running SQL statements● UPDATE: A job that updates data.● DELETE: A job that deletes a SQL job.● DATA_MIGRATION: A job that migrates data.● RESTART_QUEUE: A job that restarts a queue.● SCALE_QUEUE: A job that changes queue specifications, including scale-out and scale-in.
Status	Job status. Possible values are as follows: <ul style="list-style-type: none">● Submitting● Running● Finished● Canceled● Failed● Scaling
Query	SQL statements for operations such as exporting and creating tables You can click  to copy the query statement.
Duration	Running duration of a job
Created	Time when a job is created. Jobs can be displayed in ascending or descending order of the job creation time.

Parameter	Description
Operation	<ul style="list-style-type: none"> ● Edit: Edit the job. ● Cancel <ul style="list-style-type: none"> - You can terminate a job only when the job is in Submitting or Running status. - A job whose status is Finished, Failed, or Canceled cannot be terminated. - If the Cancel button is gray, you are not allowed to perform this operation. ● Re-execute: Execute the job again. ● SparkUI: Display the Spark job execution page. <p>NOTE</p> <ul style="list-style-type: none"> ● When you execute a job on a created queue, the cluster is restarted. It takes about 10 minutes. If you click SparkUI before the cluster is created, an empty projectID will be cached. The SparkUI page cannot be displayed. You are advised to use a dedicated queue so that the cluster will not be released. Alternatively, wait for a while after the job is submitted (the cluster is created), and then check SparkUI. ● Currently, only the latest 100 job information records are displayed on the SparkUI of DLI. <ul style="list-style-type: none"> ● In addition to the preceding operations, the following operations are available for QUERY jobs and asynchronous DDL jobs. <ul style="list-style-type: none"> - Submit Download Request: Download the results of asynchronous DDL and QUERY statements to a local directory. For details, see Submitting Download Request. - View Result: View the job running result. - Export Result: Export the job running result to the created OBS bucket. For details, see Exporting Query Results. ● In addition to the preceding operations, the EXPORT job also includes the following operations: <ul style="list-style-type: none"> - Download ● View Log: Save job logs to the temporary OBS bucket created by DLI. <p>NOTE</p> <p>The View Log button is not available for synchronization jobs and jobs running on the default queue.</p>

Searching for a Job

On the **SQL Jobs** page, you can search jobs with any of the following operations.

- Select a queue name.

- Select an engine.
- Set the date range.
- Enter a username, statement, tag, or job ID.
- Select the creation time in ascending or descending order.
- Select a job type.
- Select a job status.
- Select the job execution duration in ascending or descending order.

Viewing Job Details

On the **SQL Jobs** page, you can click  in front of a job record to view details about the job.

Job details vary with job types. The job details vary depending on the job types, status, and configuration options. The following describes how to load data, create a table, and select a job. For details about other job types, see the information on the management console.

- **Load data** (job type: IMPORT) include the following information: queue, job ID, username, type, status, execution statement, running duration, creation time, end time, parameter settings, label, number of results, scanned data, number of scanned data, number of error records, storage path, data format, database, table, table header, separator, reference character, escape character, date format, timestamp format, total CPU used, and output bytes.
- **Create table** (job type: DDL) include the following information: queue, job ID, username, type, status, execution statement, running duration, creation time, end time, parameter settings, tags, number of results, scanned data, and database.
- **Select** (job type: QUERY) include the following information: queue, job ID, username, type, status, execution statement, running duration, creation time, end time, parameter setting, label, number of results (results of successful executions can be exported), and scanned data, username, result status (results of successful tasks can be viewed. Failure causes of failed tasks are displayed), database, total CPU used, and output bytes.

NOTE

- **Total CPU Used (Core x ms)**: total CPU used during job execution.
- **Output Bytes**: number of output bytes after the job is executed.

Terminating a Job

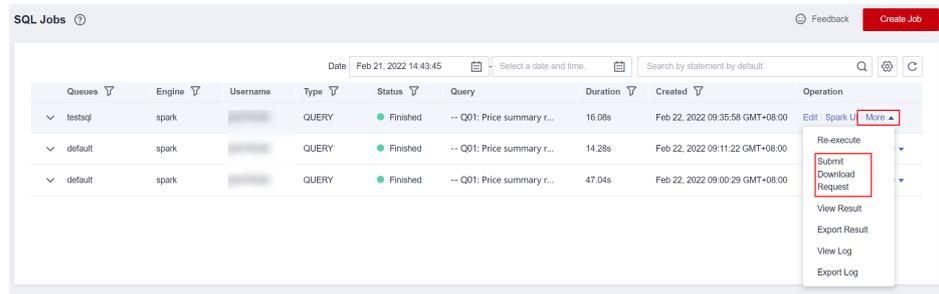
On the **SQL Jobs** page, you can click **Terminate** in the **Operation** column to stop a submitting or running job.

Submitting Download Request

Results of asynchronous DDL and QUERY statements can be downloaded to a local directory. The procedure is as follows:

1. Click **Submit Download Request** in the **Operation** column of the asynchronous DDL or QUERY job that is successfully executed. In the dialog

box that is displayed, click **OK**. After a few seconds, the **Submit Download Request** button would change to **Download**.



2. Click **Download** to download the result to your local PC.

Exporting Query Results

A maximum of 1000 records can be displayed in the query result on the console. To view more or all data, you can export the data to OBS. The procedure is as follows:

You can export results on the **SQL Jobs** page or the **SQL Editor** page.

- On the **Job Management > SQL Jobs** page, you can click **More > Export Result** in the **Operation** column to export the query result.
- After the query statements are successfully executed on the **SQL Editor** page, click  on the right of the **View Result** tab page to export the query result.

NOTE

If no column of the numeric type is displayed in the query result, the result cannot be exported.

Figure 3-1 Exporting the result

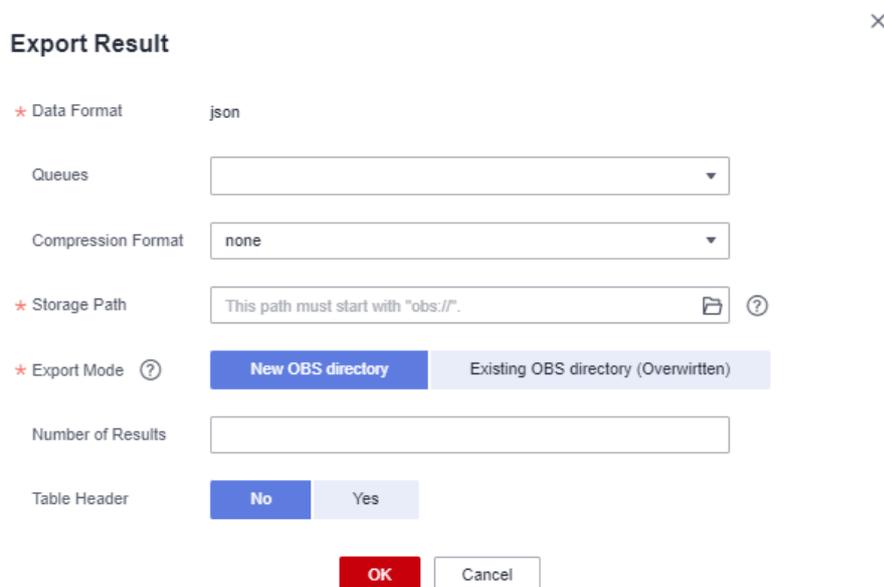


Table 3-2 Exporting parameters

Parameter	Description
Data Format	Format of the exported query result file. This parameter can be set to json or csv .
Queues	The queue where the jobs are executed. SQL jobs can be executed only in SQL queues. For details about how to create a queue, see Creating a Queue .
Compression Format	Compression format of the data to be exported. The following options are supported: <ul style="list-style-type: none">• none• bzip2• deflate• gzip
Storage Path	OBS path to store the result. NOTE <ul style="list-style-type: none">• After selecting an OBS bucket, enter a name for the folder. If the folder does not exist, it will be created in OBS.• The folder name cannot contain the special characters of \ / : * ? "< > , and cannot start or end with a dot (.
Export Mode	Mode for saving the exported query result data. <ul style="list-style-type: none">• New OBS directory: If the specified export directory exists, an error is reported and the export operation cannot be performed.• Existing OBS directory (Overwritten): If you create a file in the specified directory, the existing file will be overwritten.
Number of Results	Number of exported query results. If no value is entered or the value is 0 , all results are exported.
Table Header	Whether the data to be exported contains table headers.

3.3 Flink Job Management

3.3.1 Overview

On the Job Management page of Flink jobs, you can submit a Flink job. Currently, the following job types are supported:

- **Flink SQL** uses SQL statements to define jobs and can be submitted to any general purpose queue.
- **Flink Jar** customizes a JAR package job based on Flink APIs. It runs on dedicated queues.

- **Flink OpenSource SQL** is compatible with the SQL syntax of the community Flink 1.10 version and can run on CCE queues only.

Flink job management provides the following functions:

- [Managing Flink Job Permissions](#)
- [Creating a Flink SQL Job](#)
- [Creating a Flink Jar Job](#)
- [Editing a job](#)
- [Starting a Job](#)
- [Stopping a Job](#)
- [Deleting a Job](#)
- [Exporting a Job](#)
- [Importing a Job](#)
- [Modifying Name and Description](#)
- [Importing to a Savepoint](#)
- [Triggering a Savepoint](#)
- [Runtime Configuration](#)
- [Job Details](#)
- [Tag Management](#)

Assigning Agency Permissions

Agencies are required for DLI to execute Flink jobs. You can set the agency when logging in to the management console for the first time or go to **Global Configurations** > [Service Authorization](#) to modify the agencies.

The permissions are as follows:

- **Tenant Administrator (global)** permissions are required to access data from OBS to execute Flink jobs on DLI, for example, obtaining OBS/GaussDB(DWS) data sources, log dump (including bucket authorization), checkpointing enabling, and job import and export.

NOTE

Due to cloud service cache differences, permission setting operations require about 60 minutes to take effect.

- **DIS Administrator** permissions are required to use DIS data as the data source of DLI Flink jobs.

NOTE

Due to cloud service cache differences, permission setting operations require about 30 minutes to take effect.

- To use CloudTable data as the data source of DLI Flink jobs, **CloudTable Administrator** permissions are required.

NOTE

Due to cloud service cache differences, permission setting operations require about 3 minutes to take effect.

Flink Jobs Page

On the **Overview** page, click **Flink Jobs** to go to the Flink job management page. Alternatively, you can choose **Job Management > Flink Jobs** from the navigation pane on the left. The page displays all Flink jobs. If there are a large number of jobs, they will be displayed on multiple pages. DLI allows you to view jobs in all statuses.

Table 3-3 Job management parameters

Parameter	Description
ID	ID of a submitted Flink job, which is generated by the system by default.
Name	Name of the submitted Flink job.
Type	Type of the submitted Flink job. Including: <ul style="list-style-type: none">• Flink SQL: Flink SQL jobs• Flink Jar: Flink Jar jobs• Flink OpenSource SQL: Flink OpenSource SQL jobs
Status	Job statuses, including: <ul style="list-style-type: none">• Draft• Submitting• Submission failed• Running: The billing starts. After the job is submitted, a normal result is returned.• Running exception: The billing stops. The job stops running due to an exception.• Downloading• Idle• Stopping• Stopped• Stopping failed• Creating the savepoint• Stopped due to arrears: The billing ends. The job is stopped because the user account is in arrears.• Restoring (recharged jobs): The account in arrears is recharged, and the job is being restored.• Completed
Description	Description of the submitted Flink job.
Username	Name of the user who submits a job.
Created	Time when a job is created.
Started	Time when a Flink job starts to run.

Parameter	Description
Duration	Time consumed by job running.
Operation	<ul style="list-style-type: none">● Edit: Edit a created job. For details, see Editing a Job.● Start: Start and run a job. For details, see Starting a Job.● More<ul style="list-style-type: none">– FlinkUI: After you click this button, the Flink job execution page is displayed. NOTE When you execute a job on a created queue, the cluster is restarted. It takes about 10 minutes. If you click FlinkUI before the cluster is created, an empty projectID will be cached. The FlinkUI page cannot be displayed. You are advised to use a dedicated queue so that the cluster will not be released. Alternatively, wait for a while after the job is submitted (the cluster is created), and then check FlinkUI.– Stop: Stop a Flink job. If this function is unavailable, jobs in the current status cannot be stopped.– Delete: Delete a job. NOTE A deleted job cannot be restored.– Modify Name and Description: You can modify the name and description of a job. For details, see Modifying Name and Description.– Import Savepoint: Import the data exported from the original CS job. For details, see Importing to a Savepoint.– Trigger Savepoint: You can click this button for jobs in the Running status to save the job status. For details, see Triggering a Savepoint.– Permissions: You can view the user permissions corresponding to the job and grant permissions to other users. For details, see Managing Flink Job Permissions.– Runtime Configuration: You can enable Alarm Generation upon Job Exception and Auto Restart upon Exception. For details, see Runtime Configuration.

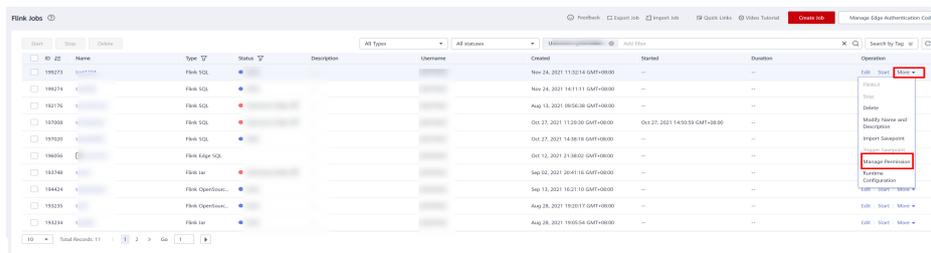
3.3.2 Managing Flink Job Permissions

Scenario

- You can isolate Flink jobs allocated to different users by setting permissions to ensure data query performance.
- The administrator and job creator have all permissions, which cannot be set or modified by other users.
- When setting job permissions for a new user, ensure that the region of the user group to which the user belongs has the **Tenant Guest** permission.

Flink Job Permission Operations

1. On the left of the DLI management console, choose **Job Management > Flink Jobs**.
2. Select the job to be configured and choose **More > Permissions** in the **Operation** column. The **User Permissions** area displays the list of users who have permissions on the job.



You can assign queue permissions to new users, modify permissions for users who have some permissions of a queue, and revoke all permissions of a user on a queue.

- Assign permissions to a new user.
 - A new user does not have permissions on the job.
 - i. Click **Grant Permission** on the right of **User Permissions** page. The **Grant Permission** dialog box is displayed.
 - ii. Specify **Username** and select corresponding permissions.
 - iii. Click **OK**.

Table 3-4 describes the related parameters.

Figure 3-2 Granting permissions

Grant Permission

* Username

Select the permissions to be granted to the user

Select all
 Get Job Details
 Start job
 Grant Permission
 Modify Jobs
 Stop job
 Revoke Permission
 Delete Jobs
 Export Job
 View Other User's Permissions

Table 3-4 Permission parameters

Parameter	Description
Username	Name of the user you want to grant permissions to. NOTE The username is the name of an existing IAM user. In addition, the user can perform authorization operations only after logging in to the Huawei Cloud platform.

Parameter	Description
Permissions to be granted to the user	<ul style="list-style-type: none"> ● Select all: All permissions are selected. ● View Job Details: This permission allows you to view the job details. ● Modify Job: This permission allows you to modify the job. ● Delete Job: This permission allows you to delete the job. ● Start Job: This permission allows you to start the job. ● Stop Job: This permission allows you to stop the job. ● Export Job: This permission allows you to export the job. ● Grant Permission: This permission allows you to grant job permissions to other users. ● Revoke Permission: This permission allows you to revoke the job permissions that other users have but cannot revoke the job creator's permissions. ● View Other User's Permissions: This permission allows you to view the job permissions of other users.

- To assign or revoke permissions of a user who has some permissions on the job, perform the following steps:
 - i. In the list under **User Permissions** for a job, select the user whose permissions need to be modified and click **Set Permission** in the **Operation** column.
 - ii. In the displayed **Set Permission** dialog box, modify the permissions of the current user. [Table 3-4](#) lists the detailed permission descriptions.
If all options under **Set Permission** are gray, you are not allowed to change permissions on this job. You can apply to the administrator, job creator, or other authorized users for job permission granting and revoking.
 - iii. Click **OK**.
- To revoke all permissions of a user on a job, perform the following steps:
In the list under **User Permissions** for a job, locate the user whose permissions need to be revoked, click **Revoke Permission** in the **Operation** column, and click **Yes**. After this operation, the user does not have any permission on the job.

Flink Job Permissions

- **View Job Details**
 - Tenants and the admin user can view and operate all jobs.

- Subusers and users with the read-only permission can only view their own jobs.

 **NOTE**

If another user grants any permission other than the job viewing permission to a subuser, the job is displayed in the job list, but the details cannot be viewed by the subuser.

- **Start Job**
 - To use a dedicated queue, you must have the permission to submit and start jobs.
 - To use a shared queue, you only need to have the permission to start jobs.
- **Stop Job**
 - To use a dedicated queue, you must have the permission to stop jobs and queues.
 - To use a shared queue, you only need to have the permission to stop jobs.
- **Delete Job**
 - If a job can be deleted, you can delete the job if you were granted this permission.
 - If a job cannot be deleted, the system stops the job before you delete it. For details about how to stop a job, see [Stop Job](#). In addition, you must have the permission to delete the job.
- **Create Job**
 - By default, sub-users cannot create jobs.
 - To create a job, you must have this permission. Currently, only the admin user has the permission to create jobs. In addition, the user must have the permission of the related package group or package used by the job.
- **Modify Job**

When modifying a job, you need to have the permission to update the job and the permission to the package group or package used by the job belongs.

3.3.3 Preparing Flink Job Data

To create a Flink job, you need to enter the data source and data output channel, that is, source and sink. To use another service as the source or sink stream, you need to apply for the service first.

Flink jobs support the following data sources and output channels:

- DIS as the data input and output channel

To use DIS as the data source and output channel, you need to enable DIS first.

For details about how to create a DIS stream, see **Creating a DIS Stream** in the *Data Ingestion Service User Guide*.

After applying for a DIS stream, you can upload local data to DIS to provide data sources for Flink jobs in real time. For details, see **Sending Data to DIS** in the *Data Ingestion Service User Guide*.

An example is provided as follows:

1,lilei,bmw320i,28
2,hanmeimei,audia4,27

- **OBS as the data source**
After you enable OBS, upload local files to OBS using the Internet. For detailed operations, see **Uploading a File** in the *Object Storage Service Console Operation Guide*.
- **RDS as the output channel**
To use RDS as the output channel, create an RDS instance. For details, see **Creating a DB Instance** in the *Relational Database Service User Guide*.
- **SMN as the output channel**
To use SMN as the output channel, create an SMN topic to obtain the URN resource ID and then add topic subscription. For detailed operations, see **Getting Started** in the *Simple Message Notification User Guide*.
- **Kafka as the data input and output channel**
If Kafka serves as both the source and sink streams, create an enhanced datasource connection between Flink jobs and Kafka. For details, see **Enhanced Datasource Connections**.
If the port of the Kafka server is listened on by the host name, you need to add the mapping between the host name and IP address of the Kafka Broker node to the datasource connection.
- **CloudTable as the data input and output channel**
To use CloudTable as the data input and output channel, create a cluster in CloudTable and obtain the cluster ID.
- **CSS as the output channel**
To use CSS as the data output channel, create a cluster in CSS and obtain the cluster's private network address. For details, see **Getting Started** in the *Cloud Search Service User Guide*.
- **DCS as the output channel**
To use DCS as the output channel, create a Redis cache instance in DCS and obtain the address used for Flink jobs to connect to the Redis instance. For detailed operations, see **Getting Started** in the *Distributed Cache Service User Guide*.

3.3.4 (Recommended) Creating a Flink OpenSource SQL Job

This section describes how to create a Flink OpenSource SQL job.

DLI Flink OpenSource SQL jobs are fully compatible with the syntax of Flink 1.10 and 1.12 provided by the community. In addition, Redis, GaussDB(DWS), and DIS data source types are added based on the community connector. For details about the syntax and restrictions of Flink SQL DDL, DML, and functions, see **Table API & SQL**.

Prerequisites

- You have prepared the data input and data output channels. For details, see **Preparing Flink Job Data**.
- Before creating a Flink OpenSource SQL job to access other external data sources, such as OpenTSDB, HBase, Kafka, GaussDB(DWS), RDS, CSS,

CloudTable, DCS Redis, and DDS MongoDB, you need to create a cross-source connection to connect the job running queue to the external data source.

- For details about the external data sources that can be accessed by Flink jobs, see [Cross-Source Analysis Development Methods](#).
- For details about how to create a datasource connection, see [Enhanced Datasource Connections](#).

On the **Resources > Queue Management** page, locate the queue you have created, and choose **More > Test Address Connectivity** in the **Operation** column to check whether the network connection between the queue and the data source is normal. For details, see [Testing Address Connectivity](#).

Precautions

Before creating jobs and submitting tasks, you are advised to enable CTS to record DLI operations for queries, audits, and tracking. [DLI Operations That Can Be Recorded by CTS](#) lists DLI operations that can be recorded by CTS.

For details about how to enable CTS and view trace details, see [Cloud Trace Service Getting Started](#).

Creating a Flink OpenSource SQL Job

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 In the upper right corner of the **Flink Jobs** page, click **Create Job**.

Figure 3-3 Creating a Flink OpenSource SQL job

Create Job
×

Type

* Name

Description

Template Name

Tags It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#) ↻

To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

Step 3 Specify job parameters.

Table 3-5 Job configuration information

Parameter	Description
Type	Set Type to Flink OpenSource SQL . You will need to start jobs by compiling SQL statements.
Name	Name of a job. Enter 1 to 57 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. NOTE The job name must be globally unique.
Description	Description of a job. It can be up to 512 characters long.
Template Name	You can select a sample template or a custom job template. For details about templates, see Managing Flink Templates .

Parameter	Description
Tag	<p>Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS).</p> <p>For details, see Tag Management Service User Guide.</p> <p>NOTE</p> <ul style="list-style-type: none">• A maximum of 20 tags can be added.• Only one tag value can be added to a tag key.• The key name in each resource must be unique.• Tag key: Enter a tag key name in the text box. <p>NOTE</p> <p>A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed, but the value cannot start or end with a space or start with <code>_sys_</code>.</p> <ul style="list-style-type: none">• Tag value: Enter a tag value in the text box. <p>NOTE</p> <p>A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed. The value cannot start or end with a space.</p>

Step 4 Click **OK** to enter the editing page.

Step 5 Edit an OpenSource SQL job.

Enter details SQL statements in the SQL statement edit area. For details about SQL statements, see the *Data Lake Insight Flink OpenSource SQL Syntax Reference*.

Step 6 Click **Check Semantics**.

- You can **Start** a job only after the semantic verification is successful.
- If verification is successful, the message "The SQL semantic verification is complete. No error." will be displayed.
- If verification fails, a red "X" mark will be displayed in front of each SQL statement that produced an error. You can move the cursor to the "X" mark to view error details and change the SQL statement as prompted.

Step 7 Set job running parameters.

Figure 3-4 Setting running parameters for Flink OpenSource SQL

The screenshot shows a configuration panel for Flink OpenSource SQL with the following settings:

- Queue:** Shared queue (Free for O...)
- CUs:** 2
- Job Manager CUs:** 1
- Parallelism:** 1
- Task Manager Configu...:**
- OBS Bucket:** dli-
- Save Job Log:**
- Alarm Generation upo...:**
- Enable Checkpointing:**
- Checkpoint Interval:** 30 s
- Checkpoint Mode:** Exactly once
- Auto Restart upon Exc...:**
- Idle State Retention Time:** 1 h

Table 3-6 Running parameters

Parameter	Description
Queue	<p>A shared queue is selected by default. Select a CCE queue with dedicated resources and configure the following parameters:</p> <p>UDF Jar: You can customize a UDF Jar file. Before selecting a JAR file to be inserted, upload the corresponding JAR file to the OBS bucket and choose Data Management > Package Management to create a package. For details, see Creating a Package.</p> <p>In SQL, you can call a user-defined function that is inserted into a JAR file.</p> <p>NOTE</p> <p>During job creation, a sub-user can only select a queue that has been allocated to the user.</p> <p>If the remaining capacity of the selected queue cannot meet the job requirements, the system automatically scales up the capacity and you will be billed based on the increased capacity. When a queue is idle, the system automatically scales in the queue.</p>

Parameter	Description
CUs	<p>Sum of the number of compute units and job manager CUs of DLI. CU is also the billing unit of DLI. One CU equals 1 vCPU and 4 GB.</p> <p>The value is the number of CUs required for job running and cannot exceed the number of CUs in the bound queue.</p>
Job Manager CUs	Number of CUs of the management unit.
Parallelism	<p>Number of Flink OpenSource SQL jobs that run at the same time</p> <p>NOTE This value cannot be greater than four times the compute units (number of CUs minus the number of job manager CUs).</p>
Task Manager Configuration	<p>Whether to set Task Manager resource parameters.</p> <p>If this option is selected, you need to set the following parameters:</p> <ul style="list-style-type: none">• CU(s) per TM: Number of resources occupied by each Task Manager.• Slot(s) per TM: Number of slots contained in each Task Manager.
OBS Bucket	OBS bucket to store job logs and checkpoint information. If the selected OBS bucket is not authorized, click Authorize .
Save Job Log	<p>Whether to save the job running logs to OBS. The logs are saved in the following path: <i>Bucket name/jobs/logs/Directory starting with the job ID</i>.</p> <p>CAUTION You are advised to configure this parameter. Otherwise, no run log is generated after the job is executed. If the job fails, the run log cannot be obtained for fault locating.</p> <p>If this option is selected, you need to set the following parameters:</p> <p>OBS Bucket: Select an OBS bucket to store user job logs. If the selected OBS bucket is not authorized, click Authorize.</p> <p>NOTE If both Enable Checkpointing and Save Job Log are selected, you only need to authorize OBS once.</p>
Alarm Generation upon Job Exception	<p>Whether to report job exceptions, for example, abnormal job running or exceptions due to an insufficient balance, to users via SMS or email.</p> <p>If this option is selected, you need to set the following parameters:</p> <p>SMN Topic Select a user-defined SMN topic. For details about how to customize SMN topics, see Creating a Topic in the <i>Simple Message Notification User Guide</i>.</p>

Parameter	Description
Enable Checkpointing	<p>Whether to enable job snapshots. If this function is enabled, jobs can be restored based on the checkpoints.</p> <p>If this option is selected, you need to set the following parameters:</p> <ul style="list-style-type: none"> • Checkpoint Interval (s) indicates the interval for creating checkpoints. The value ranges from 1 to 999999, and the default value is 30. • Checkpoint Mode can be set to either of the following values: <ul style="list-style-type: none"> – At least once: Events are processed at least once. – Exactly once: Events are processed only once. • OBS Bucket: Select an OBS bucket to store your checkpoints. If the selected OBS bucket is not authorized, click Authorize. The checkpoint path is <i>Bucket name/jobs/checkpoint/Directory starting with the job ID</i>. <p>NOTE If both Enable Checkpointing and Save Job Log are selected, you only need to authorize OBS once.</p>
Auto Restart upon Exception	<p>Whether to enable automatic restart. If this function is enabled, any job that has become abnormal will be automatically restarted.</p> <p>If this option is selected, you need to set the following parameters:</p> <ul style="list-style-type: none"> • Max. Retry Attempts: maximum number of retry times upon an exception. The unit is Times/hour. <ul style="list-style-type: none"> – Unlimited: The number of retries is unlimited. – Limited: The number of retries is user-defined. • Restore Job from Checkpoint: This parameter is available only when Enable Checkpointing is selected.
Idle State Retention Time	<p>How long the state of a key is retained without being updated before it is removed in GroupBy or Window. The default value is 1 hour.</p>
Dirty Data Policy	<p>Policy for processing dirty data. The following policies are supported: Ignore, Trigger a job exception, and Save.</p> <p>If you set this field to Save, the Dirty Data Dump Address must be set. Click the address box to select the OBS path for storing dirty data.</p>

Step 8 (Optional) Set the runtime configuration as required.

Figure 3-5 Runtime configuration

Step 9 Click **Save**.

Step 10 Click **Start**. On the displayed **Start Flink Jobs** page, confirm the job specifications and the price, and click **Start Now** to start the job.

After the job is started, the system automatically switches to the **Flink Jobs** page, and the created job is displayed in the job list. You can view the job status in the **Status** column. After a job is successfully submitted, the job status will change from **Submitting** to **Running**. After the execution is complete, the message **Completed** is displayed.

If the job status is **Submission failed** or **Running exception**, the job submission failed or the job did not execute successfully. In this case, you can move the cursor over the status icon in the **Status** column of the job list to view the error details. You can click  to copy these details. After handling the fault based on the provided information, resubmit the job.

NOTE

Other buttons are as follows:

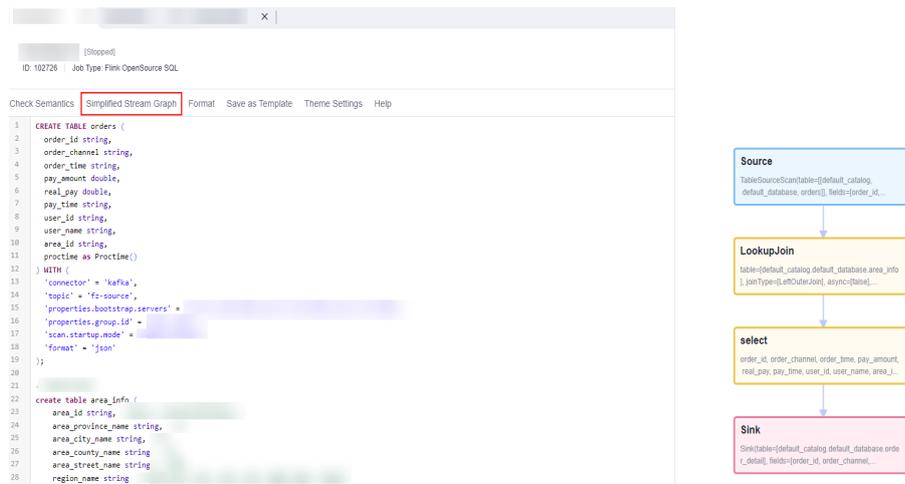
- **Save As:** Save the created job as a new job.
- **Static Stream Graph:** Provide the static concurrency estimation function and stream graph display function. See [Figure 3-7](#).
- **Simplified Stream Graph:** Display the data processing flow from the source to the sink. See [Figure 3-6](#).
- **Format:** Format the SQL statements in the editing box.
- **Set as Template:** Set the created SQL statements as a job template.
- **Theme Settings:** Set the theme related parameters, including **Font Size**, **Wrap**, and **Page Style**.
- **Help:** Redirect to the help center to provide you with the SQL syntax for stream jobs.

----End

Simplified Stream Graph

On the OpenSource SQL job editing page, click **Simplified Stream Graph**.

Figure 3-6 Simplified stream graph



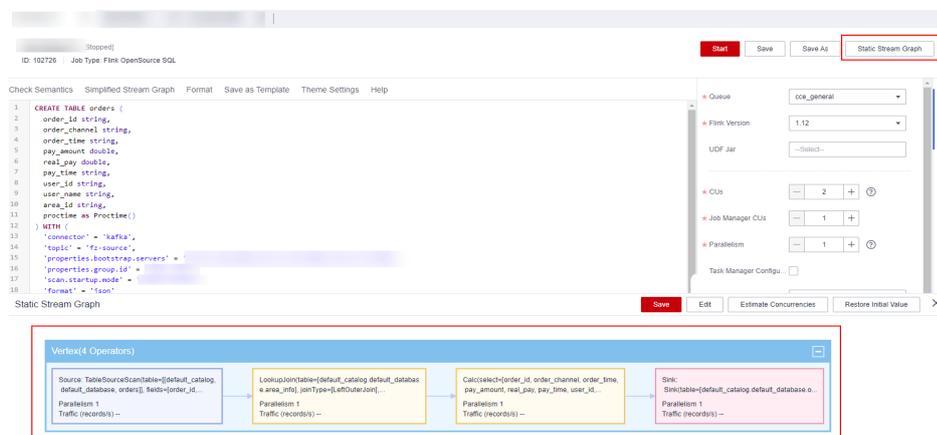
Static Stream Graph

On the OpenSource SQL job editing page, click **Static Stream Graph**.

The **Static Stream Graph** page also allows you to:

- Estimate concurrencies. Click **Estimate Concurrencies** on the **Static Stream Graph** page to estimate concurrencies. Click **Restore Initial Value** to restore the initial value after concurrency estimation.
- Zoom in or out the page.
- Expand or merge operator chains.
- You can edit **Parallelism**, **Output rate**, and **Rate factor**.
 - **Parallelism:** indicates the number of concurrent tasks.
 - **Output rate:** indicates the data traffic of an operator. The unit is piece/s.
 - **Rate factor:** indicates the retention rate after data is processed by operators. Rate factor = Data output volume of an operator/Data input volume of the operator (Unit: %)

Figure 3-7 Static stream graph



3.3.5 Creating a Flink SQL Job

This section describes how to create a Flink SQL job. You can use Flink SQLs to develop jobs to meet your service requirements. Using SQL statements simplifies logic implementation. You can edit Flink SQL statements for your job in the DLI SQL editor. This section describes how to use the SQL editor to write Flink SQL statements.

DLI Flink OpenSource SQL jobs are fully compatible with the syntax of Flink 1.10 and 1.12 provided by the community. In addition, Redis, GaussDB(DWS), and DIS data source types are added based on the community connector.

Prerequisites

- You have prepared the data input and data output channels. For details, see [Preparing Flink Job Data](#).
- When you use a Flink SQL job to access other external data sources, such as OpenTSDB, HBase, Kafka, DWS, RDS, CSS, CloudTable, DCS Redis, and DDS MongoDB, you need to create a cross-source connection to connect the job running queue to the external data source.
 - For details about the external data sources that can be accessed by Flink jobs, see [Cross-Source Analysis Development Methods](#).
 - For details about how to create a datasource connection, see [Enhanced Datasource Connections](#). After a datasource connection is created, you can choose **More > Test Address Connectivity** in the **Operation** column on the **Queue Management** page to check whether the network connection between the queue and the external data source is normal. For details, see [Testing Address Connectivity](#).

Precautions

Before creating jobs and submitting tasks, you are advised to enable CTS to record operations associated with DLI for later query, audit, and backtrack operations. DLI operations that can be recorded by CTS. For details, see [DLI Operations That Can Be Recorded by CTS](#).

For details about how to enable CTS and view trace details, see [Cloud Trace Service Getting Started](#).

Creating a Flink SQL Job

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** In the upper right corner of the **Flink Jobs** page, click **Create Job**.

Figure 3-8 Creating a Flink SQL job

Create Job
×

Type

★ Name

Description

Template Name

Tags It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#)

To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

Step 3 Specify job parameters.

Table 3-7 Job configuration information

Parameter	Description
Type	Set Type to Flink SQL . You will need to rewrite SQL statements to start the job.
Name	Name of a job. Enter 1 to 57 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. NOTE The job name must be globally unique.
Description	Description of a job. It can contain a maximum of 512 characters.
Template Name	You can select a sample template or a custom job template. For details about templates, see Managing Flink Templates .

Parameter	Description
Tag	<p>Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS).</p> <p>For details, see Tag Management Service User Guide.</p> <p>NOTE</p> <ul style="list-style-type: none">• A maximum of 20 tags can be added.• Only one tag value can be added to a tag key.• The key name in each resource must be unique.• Tag key: Enter a tag key name in the text box. <p>NOTE</p> <p>A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed, but the value cannot start or end with a space or start with <code>_sys_</code>.</p> <ul style="list-style-type: none">• Tag value: Enter a tag value in the text box. <p>NOTE</p> <p>A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed. The value cannot start or end with a space.</p>

Step 4 Click **OK** to enter the editing page.

Step 5 Edit a Flink SQL job.

Enter details SQL statements in the statement editing area. For details about SQL syntax, see the Data Lake Insight SQL Syntax Reference.

Step 6 Click **Check Semantics**.

- You can **Debug** or **Start** a job only after the semantic verification is successful.
- If verification is successful, the message "The SQL semantic verification is complete. No error." will be displayed.
- If verification fails, a red "X" mark will be displayed in front of each SQL statement that produced an error. You can move the cursor to the "X" mark to view error details and change the SQL statement as prompted.

Step 7 Set job running parameters.

Table 3-8 Running parameters

Parameter	Description
Queue	<p>A shared queue is selected by default. You can select a custom queue as needed.</p> <p>NOTE</p> <ul style="list-style-type: none">• During job creation, a sub-user can only select a queue that has been allocated to the user.• If the remaining capacity of the selected queue cannot meet the job requirements, the system automatically scales up the capacity and you will be billed based on the increased capacity. When a queue is idle, the system automatically scales in the queue.
UDF Jar	<p>If you selected custom queues, you need to configure this parameter.</p> <p>You can customize a UDF Jar file. Before you select a JAR file, upload the corresponding JAR package to the OBS bucket and choose Data Management > Package Management to create a package. For details, see Creating a Package.</p> <p>In SQL, you can call a user-defined function that is inserted into a JAR file.</p>
CUs	<p>Sum of the number of compute units and job manager CUs of DLI. CU is also the billing unit of DLI. One CU equals 1 vCPU and 4 GB.</p> <p>The configured number of CUs is the number of CUs required for job running and cannot exceed the number of CUs in the bound queue.</p>
Job Manager CUs	Number of CUs of the management unit.
Parallelism	<p>Number of Flink SQL jobs that run at the same time Properly increasing the number of parallel threads improves the overall computing capability of the job. However, the switchover overhead caused by the increase of threads must be considered.</p> <p>NOTE</p> <ul style="list-style-type: none">• This value cannot be greater than four times the compute units (number of CUs minus the number of job manager CUs).• The priority of the number of parallel tasks on this page is lower than that set in the code.
Task Manager Configuration	<p>Whether to set Task Manager resource parameters</p> <p>If this option is selected, you need to set the following parameters:</p> <ul style="list-style-type: none">• CU(s) per TM: Number of resources occupied by each Task Manager.• Slot(s) per TM: Number of slots contained in each Task Manager.

Parameter	Description
OBS Bucket	<p>OBS bucket to store job logs and checkpoint information. If the selected OBS bucket is not authorized, click Authorize.</p> <p>NOTE If Enable Checkpointing and Save Job Log are both selected, you only need to authorize OBS once.</p>
Save Job Log	<p>Whether to save the job running logs to OBS The logs are saved in the following path: <i>Bucket name/jobs/logs/Directory starting with the job ID</i>. To go to this path, go to the job list and click the job name. On the Run Log tab page, click the provided OBS link.</p> <p>CAUTION You are advised to select this parameter. Otherwise, no run log is generated after the job is executed. If the job is abnormal, the run log cannot be obtained for fault locating.</p> <p>If this option is selected, you need to set the following parameters:</p> <p>OBS Bucket: Select an OBS bucket to store user job logs. If the selected OBS bucket is not authorized, click Authorize.</p> <p>NOTE If Enable Checkpointing and Save Job Log are both selected, you only need to authorize OBS once.</p>
Alarm Generation upon Job Exception	<p>Whether to report job exceptions, for example, abnormal job running or exceptions due to an insufficient balance, to users via SMS or email</p> <p>If this option is selected, you need to set the following parameters:</p> <p>SMN Topic Select a user-defined SMN topic. For details about how to create a custom SMN topic, see "Creating a Topic" in <i>Simple Message Notification User Guide</i>.</p>

Parameter	Description
Enable Checkpointing	<p>Whether to enable job snapshots. If this function is enabled, jobs can be restored based on the checkpoints.</p> <p>If this option is selected, you need to set the following parameters:</p> <ul style="list-style-type: none">• Checkpoint Interval indicates the interval for creating checkpoints. The value ranges from 1 to 999999, and the default value is 30.• Checkpoint Mode can be set to either of the following values:<ul style="list-style-type: none">– At least once: Events are processed at least once.– Exactly once: Events are processed only once.• OBS Bucket: Select an OBS bucket to store your checkpoints. If the selected OBS bucket is not authorized, click Authorize. The checkpoint path is <i>Bucket name/jobs/checkpoint/Directory starting with the job ID</i>. <p>NOTE If Enable Checkpointing and Save Job Log are both selected, you only need to authorize OBS once.</p>
Auto Restart upon Exception	<p>Whether to enable automatic restart. If this function is enabled, any job that has become abnormal will be automatically restarted.</p> <p>If this option is selected, you need to set the following parameters:</p> <ul style="list-style-type: none">• Max. Retry Attempts: maximum number of retry times upon an exception. The unit is times/hour.<ul style="list-style-type: none">– Unlimited: The number of retries is unlimited.– Limited: The number of retries is user-defined.• Restore Job from Checkpoint: This parameter is available only when Enable Checkpointing is selected.
Idle State Retention Time	<p>Defines for how long the state of a key is retained without being updated before it is removed in GroupBy or Window. The default value is 1 hour.</p>
Dirty Data Policy	<p>Select a policy for processing dirty data. The following policies are supported: Ignore, Trigger a job exception, and Save.</p> <p>NOTE Save indicates that the dirty data is stored to the OBS bucket selected above.</p>
Dirty Data Dump Address	<p>Set this parameter when Dirty Data Policy is set to Save. Click the address box to select the OBS path for storing dirty data.</p>

Step 8 (Optional) Set the runtime configuration as required.

Step 9 Click **Save**.

Step 10 Click **Start**. On the displayed **Start Flink Jobs** page, confirm the job specifications and the price, and click **Start Now** to start the job.

After the job is started, the system automatically switches to the **Flink Jobs** page, and the created job is displayed in the job list. You can view the job status in the **Status** column. After a job is successfully submitted, the job status will change from **Submitting** to **Running**. After the execution is complete, the message **Completed** is displayed.

If the job status is **Submission failed** or **Running exception**, the job submission failed or the job did not execute successfully. In this case, you can move the cursor over the status icon in the **Status** column of the job list to view the error details. You can click  to copy error information. After handling the fault based on the provided information, resubmit the job.

NOTE

Other available buttons are as follows:

- **Save As:** Save the created job as a new job.
- **Format:** Format the SQL statements in the editing box.
- **Set as Template:** Set the created SQL statements as a job template.
- **Theme Settings:** Set the theme related parameters, including **Font Size**, **Wrap**, and **Page Style**.
- **Help:** Redirect to the help center to provide you with the SQL syntax for stream jobs.

----End

3.3.6 Creating a Flink Jar Job

This section describes how to create a Flink Jar job. You can perform secondary development based on Flink APIs, build your own JAR file, and submit the JAR file to DLI queues. DLI is fully compatible with open-source community APIs. To create a custom Flink job, you need to compile and build application JAR files. You must have a certain understanding of Flink secondary development and have high requirements related to stream computing complexity.

Prerequisites

- Ensure that a dedicated queue has been created. To create a dedicated queue, select **Dedicated Resource Mode** when you choose the type of a queue during purchase.
- When creating a Flink Jar job to access other external data sources, such as OpenTSDB, HBase, Kafka, GaussDB(DWS), RDS, CSS, CloudTable, DCS Redis, and DDS MongoDB, you need to create a cross-source connection to connect the job running queue to the external data source.
 - For details about the external data sources that can be accessed by Flink jobs, see [Cross-Source Analysis Development Methods](#).
 - For details about how to create a datasource connection, see [Enhanced Datasource Connections](#).

On the **Resources > Queue Management** page, locate the queue you have created, and choose **More > Test Address Connectivity** in the

Operation column to check whether the network connection between the queue and the data source is normal. For details, see [Testing Address Connectivity](#).

- When running a Flink Jar job, you need to build the secondary development application code into a Jar package and upload the JAR package to the created OBS bucket. Choose **Data Management > Package Management** to create a package. For details, see [Creating a Package](#).

 **NOTE**

DLI does not support the download function. If you need to modify the uploaded data file, please edit the local file and upload it again.

- Flink dependencies have been built in the DLI server and security hardening has been performed based on the open-source community version. To avoid dependency package compatibility issues or log output and dump issues, be careful to exclude the following files when packaging:
 - Built-in dependencies (or set the package dependency scope to **provided** in Maven or SBT)
 - Log configuration files (example, **log4j.properties/logback.xml**)
 - JAR package for log output implementation (example, **log4j**).

Precautions

Before creating jobs and submitting tasks, you are advised to enable CTS to record operations associated with DLI for later query, audit, and backtrack operations. To view the DLI operations that can be recorded by CTS, see [DLI Operations That Can Be Recorded by CTS](#).

For details about how to enable CTS and view trace details, see the [Cloud Trace Service Getting Started](#).

Creating a Flink Jar Job

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** In the upper right corner of the **Flink Jobs** page, click **Create Job**.
- Step 3** Specify job parameters.

Table 3-9 Job configuration information

Parameter	Description
Type	Select Flink Jar .
Name	Name of a job. Enter 1 to 57 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. NOTE The job name must be globally unique.

Parameter	Description
Description	Description of a job. It can be up to 512 characters long.
Tag	<p>Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS).</p> <p>For details, see Tag Management Service User Guide.</p> <p>NOTE</p> <ul style="list-style-type: none"> • A maximum of 20 tags can be added. • Only one tag value can be added to a tag key. • The key name in each resource must be unique. • Tag key: Enter a tag key name in the text box. <p>NOTE</p> <p>A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed, but the value cannot start or end with a space or start with <code>_sys_</code>.</p> <ul style="list-style-type: none"> • Tag value: Enter a tag value in the text box. <p>NOTE</p> <p>A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed. The value cannot start or end with a space.</p>

Step 4 Click **OK** to enter the editing page.

Step 5 Select a queue. Flink Jar jobs can run only on dedicated queues.

 **NOTE**

- A Flink Jar job can run only on a pre-created dedicated queue.
- If no dedicated queue is available in the **Queue** drop-down list, create a dedicated queue and bind it to the current user.

Step 6 Configuring Flink Jar Job parameters

Table 3-10 Parameter description

Name	Description
Queue	A shared queue is selected by default. You can select a custom queue as needed.

Name	Description
Application	<p>User-defined package. Before selecting a JAR file to be inserted, upload the corresponding JAR file to the OBS bucket and choose Data Management > Package Management to create a package. For details, see Creating a Package.</p> <p>For details about built-in dependency packages, see Built-in Dependencies.</p>
Main Class	<p>The name of the JAR package to be loaded, for example, KafkaMessageStreaming.</p> <ul style="list-style-type: none"> • Default: Specified based on the Manifest file in the JAR package. • Manually assign: You must enter the class name and confirm the class arguments (separate arguments with spaces). <p>NOTE When a class belongs to a package, the main class path must contain the complete package path, for example, packageName.KafkaMessageStreaming.</p>
Class Arguments	<p>List of arguments of a specified class. The arguments are separated by spaces.</p> <p>Flink parameters support replacement of non-sensitive global variables. For example, if you add the global variable windowSize in Global Configuration > Global Variables, you can add the -windowSize {{windowSize}} parameter for the Flink Jar job.</p>
JAR Package Dependencies	<p>Select a user-defined package dependency. The dependent program packages are stored in the classpath directory of the cluster.</p> <p>Before selecting a JAR file to be inserted, upload the corresponding JAR file to the OBS bucket and choose Data Management > Package Management to create a package. Select JAR as the package type. For details, see Creating a Package.</p> <p>For details about built-in dependency packages, see Built-in Dependencies.</p>
Other Dependencies	<p>User-defined dependency files. Other dependency files need to be referenced in the code.</p> <p>Before selecting a dependency file, upload the file to the OBS bucket and choose Data Management > Package Management to create a package. The package type is not limited. For details, see Creating a Package.</p> <p>You can add the following command to the application to access the corresponding dependency file. In the command, fileName indicates the name of the file to be accessed, and ClassName indicates the name of the class that needs to access the file.</p> <pre>ClassName.class.getClassLoader().getResource("userData/fileName")</pre>

Name	Description
Flink Version	Before selecting a Flink version, you need to select the queue to which the Flink version belongs. Currently, versions 1.10, 1.11, and 1.12 are supported.
Runtime Configuration	User-defined optimization parameters. The parameter format is key=value . Flink optimization parameters support replacement non-sensitive global variable. For example, if you create global variable phase in Global Configuration > Global Variables , optimization parameter table.optimizer.agg-phase.strategy={{phase}} can be added to the Flink Jar job.

Step 7 Configure job parameters.**Table 3-11** Parameter description

Name	Description
CUs	One CU has one vCPU and 4-GB memory. The number of CUs ranges from 2 to 10,000.
Job Manager CUs	Set the number of CUs on a management unit. The value ranges from 1 to 4. The default value is 1.
Parallelism	Maximum number of parallel operators in a job. The value ranges from 1 to 10,000. NOTE <ul style="list-style-type: none"> The value must be less than or equal to four times the number of compute units (CUs minus the number of job manager CUs). You are advised to set this parameter to a value greater than that configured in the code. Otherwise, job submission may fail.
Task Manager Configuration	Whether to set Task Manager resource parameters If this option is selected, you need to set the following parameters: <ul style="list-style-type: none"> CU(s) per TM: Number of resources occupied by each Task Manager. Slot(s) per TM: Number of slots contained in each Task Manager.
Save Job Log	Whether to save the job running logs to the OBS bucket. CAUTION You are advised to select this parameter. Otherwise, no run log is generated after the job is executed. If the job is abnormal, the run log cannot be obtained for fault locating. If this option is selected, you need to set the following parameters: OBS Bucket: Select an OBS bucket to store job logs. If the selected OBS bucket is not authorized, click Authorize .

Name	Description
Alarm Generation upon Job Exception	Whether to report job exceptions, for example, abnormal job running or exceptions due to an insufficient balance, to users via SMS or email If this option is selected, you need to set the following parameters: SMN Topic Select a user-defined SMN topic. For details about how to customize SMN topics, see Creating a Topic in the <i>Simple Message Notification User Guide</i> .
Auto Restart upon Exception	Whether to enable automatic restart. If this function is enabled, any job that has become abnormal will be automatically restarted. If this option is selected, you need to set the following parameters: <ul style="list-style-type: none">● Max. Retry Attempts: maximum number of retry times upon an exception. The unit is Times/hour.<ul style="list-style-type: none">– Unlimited: The number of retries is unlimited.– Limited: The number of retries is user-defined.● Restore Job from Checkpoint: Restore the job from the saved checkpoint. If you select this parameter, you also need to set Checkpoint Path. Checkpoint Path: Select the checkpoint saving path. The checkpoint path must be the same as that you set in the application package. Note that the checkpoint path for each job must be unique. Otherwise, the checkpoint cannot be obtained.

Step 8 Click **Save** on the upper right of the page.

Step 9 Click **Start** on the upper right side of the page. On the displayed **Start Flink Job** page, confirm the job specifications and the price, and click **Start Now** to start the job.

After the job is started, the system automatically switches to the **Flink Jobs** page, and the created job is displayed in the job list. You can view the job status in the **Status** column. After a job is successfully submitted, the job status will change from **Submitting** to **Running**. After the execution is complete, the message **Completed** is displayed.

If the job status is **Submission failed** or **Running exception**, the job submission failed or the job did not execute successfully. In this case, you can move the cursor over the status icon in the **Status** column of the job list to view the error details.

You can click  to copy these details. After handling the fault based on the provided information, resubmit the job.

NOTE

Other available buttons are as follows:

Save As: Save the created job as a new job.

----End

3.3.7 Performing Operations on a Flink Job

After a job is created, you can perform operations on the job as required.

- [Editing a Job](#)
- [Starting a Job](#)
- [Stopping a Job](#)
- [Deleting a Job](#)
- [Exporting a Job](#)
- [Importing a Job](#)
- [Modifying Name and Description](#)
- [Importing to a Savepoint](#)
- [Triggering a Savepoint](#)
- [Runtime Configuration](#)

Editing a Job

You can edit a created job, for example, by modifying the SQL statement, job name, job description, or job configurations.

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** In the row where the job you want to edit locates, click **Edit** in the **Operation** column to switch to the editing page.
- Step 3** Edit the job as required.

For details about how to edit a Flink SQL job, see Step 5 to Step 7 in [Creating a Flink SQL Job](#).

For details about how to edit a user-defined Flink job, see Step 5 to Step 7 in [Creating a Flink Jar Job](#).

----End

Starting a Job

You can start a saved or stopped job.

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** Use either of the following methods to start jobs:
- Starting a single job
Select a job and click **Start** in the **Operation** column.
Alternatively, you can select the row where the job you want to start locates and click **Start** in the upper left of the job list.
 - Batch starting jobs
Select the rows where the jobs you want to start locate and click **Start** in the upper left of the job list.

After you click **Start**, the **Start Flink Jobs** page is displayed.

Step 3 On the **Start Flink Jobs** page, confirm the job information and price. If they are correct, click **Start Now**.

After a job is started, you can view the job execution result in the **Status** column.

----End

Stopping a Job

You can stop a job in the **Running** or **Submitting** state.

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Stop a job using either of the following methods:

- Stopping a job
Locate the row that contains the job to be stopped, click **More** in the **Operation** column, and select **Stop**.
Alternatively, you can select the row where the job you want to stop locates and click **Stop** in the upper left of the job list.
- Batch stopping jobs
Locate the rows containing the jobs you want to stop and click **Stop** in the upper left of the job list.

Step 3 In the displayed **Stop Job** dialog box, click **OK** to stop the job.

Figure 3-9 Stopping a job

 Are you sure you want to stop the following 1 jobs?

Name	Status	Description
	● Running	

Trigger Savepoint 

NOTE

- Before stopping a job, you can trigger a savepoint to save the job status information. When you start the job again, you can choose whether to restore the job from the savepoint.
- If you select **Trigger savepoint**, a savepoint is created. If **Trigger savepoint** is not selected, no savepoint is created. By default, the savepoint function is disabled.
- The lifecycle of a savepoint starts when the savepoint is triggered and stops the job, and ends when the job is restarted. The savepoint is automatically deleted after the job is restarted.

When a job is being stopped, the job status is displayed in the **Status** column of the job list. The details are as follows:

- **Stopping:** indicates that the job is being stopped.
- **Stopped:** indicates that the job is stopped successfully.
- **Stop failed:** indicates that the job failed to be stopped.

----End

Deleting a Job

If you do not need to use a job, perform the following operations to delete it. A deleted job cannot be restored. Therefore, exercise caution when deleting a job.

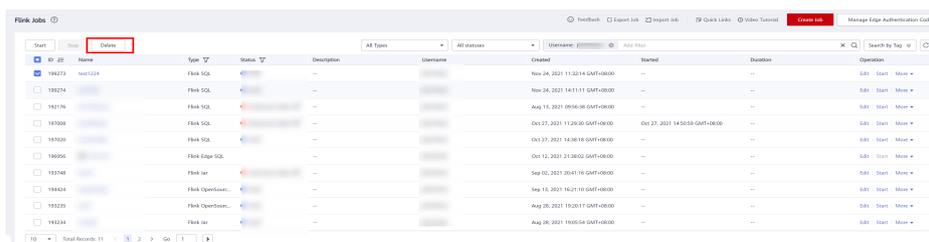
Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Perform either of the following methods to delete jobs:

- Deleting a single job

Locate the row containing the job you want to delete and click **More > Delete** in the **Operation** column.

Alternatively, you can select the row containing the job you want to delete and click **Delete** in the upper left of the job list.



- Deleting jobs in batches

Select the rows containing the jobs you want to delete and click **Delete** in the upper left of the job list.

Step 3 Click **Yes**.

----End

Exporting a Job

You can export the created Flink jobs to an OBS bucket.

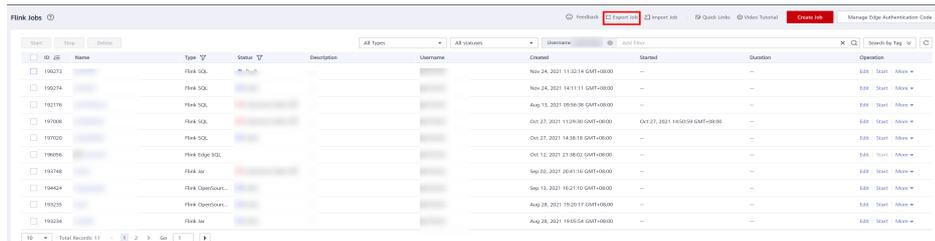
This mode is applicable to the scenario where a large number of jobs need to be created when you switch to another region, project, or user. In this case, you do not need to create a job. You only need to export the original job, log in to the system in a new region or project, or use a new user to import the job.

NOTE

When switching to another project or user, you need to grant permissions to the new project or user. For details, see [Managing Flink Job Permissions](#).

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Click **Export Job** in the upper right corner. The **Export Job** dialog box is displayed.



Step 3 Select the OBS bucket where the job is stored. Click **Next**.

Step 4 Select job information you want to export.

By default, configurations of all jobs are exported. You can enable the **Custom Export** function to export configurations of the desired jobs.

Step 5 Click **Confirm** to export the job.

----End

Importing a Job

You can import the Flink job configuration file stored in the OBS bucket to the **Flink Jobs** page of DLI.

This mode is applicable to the scenario where a large number of jobs need to be created when you switch to another region, project, or user. In this case, you do not need to create a job. You only need to export the original job, log in to the system in a new region or project, or use a new user to import the job.

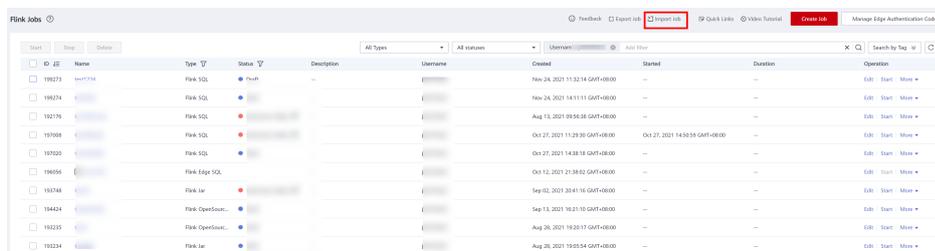
If you need to import a self-created job, you are advised to use the job creation function. For details, see [Creating a Flink SQL Job](#) and [Creating a Flink Jar Job](#).

NOTE

- When switching to another project or user, you need to grant permissions to the new project or user. For details, see [Managing Flink Job Permissions](#).
- Only jobs whose data format is the same as that of Flink jobs exported from DLI can be imported.

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Click **Import Job** in the upper right corner. The **Import Job** dialog box is displayed.



Step 3 Select the complete OBS path of the job configuration file to be imported. Click **Next**.

Step 4 Configure the same-name job policy and click next. Click **Next**.

- Select **Overwrite job of the same name**. If the name of the job to be imported already exists, the existing job configuration will be overwritten and the job status switches to **Draft**.
- If **Overwrite job of the same name** is not selected and the name of the job to be imported already exists, the job will not be imported.

Step 5 Ensure that **Config File** and **Overwrite Same-Name Job** are correctly configured. Click **Confirm** to import the job.

----End

Modifying Name and Description

You can change the job name and description as required.

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** In the **Operation** column of the job whose name and description need to be modified, choose **More > Modify Name and Description**. The **Modify Name and Description** dialog box is displayed. Change the name or modify the description of a job.

Figure 3-10 Modifying name and description

Modify Name and Description

The dialog box titled "Modify Name and Description" contains the following elements:

- A label "★ Name" followed by a text input field containing the text "testfordocument".
- A label "Description" followed by a larger, empty text area.
- At the bottom, there are two buttons: a red "OK" button and a white "Cancel" button with a grey border.

Step 3 Click **OK**.

----End

Triggering a Savepoint

When you need to stop a job, you can create a savepoint to save the job status information. In this case, when you restart the job, you can choose to restore the job from the latest savepoint.

 **NOTE**

- You can click **Trigger Savepoint** for jobs in the **Running** status to save the job status.
- The lifecycle of a savepoint starts when the savepoint is triggered and stops the job, and ends when the job is restarted. The savepoint is automatically deleted after the job is restarted.

Importing to a Savepoint

You can import a savepoint to restore the job status. For details about the savepoint, see [Checkpointing](#) at the official website of Flink.

You need to select the OBS path of the save point.

Runtime Configuration

You can select **Runtime Configuration** to configure job exception alarms and restart options.

 **NOTE**

Flink SQL jobs, Flink OpenSource SQL jobs, and Flink Jar jobs are supported.

1. In the **Operation** column of the Flink job, choose **More > Runtime Configuration**.
2. In the **Runtime Configuration** dialog box, set the following parameters:

Figure 3-11 Runtime configuration

✕

Runtime Configuration

Name testflinkjar

Alarm Generation upon Job Exception

* SMN Topic DLI_fink_info ▼
[Configure Topic](#)

Auto Restart upon Exception

Max. Retry Attempts
 Unlimited
 Limited

Restore Job from Checkpoint

* Checkpoint Path dli-test-obs01/ 📁

The checkpoint path must be the same as that you set in the application package. Note that the checkpoint path for each job must be unique. Otherwise, the checkpoint cannot be obtained.

Table 3-12 Running parameters

Parameter	Description
Name	Job name.
Alarm Generation upon Job Exception	Whether to report job exceptions, for example, abnormal job running or exceptions due to an insufficient balance, to users via SMS or email. If this option is selected, you need to set the following parameters: SMN Topic Select a user-defined SMN topic. For details about how to customize SMN topics, see Creating a Topic in the <i>Simple Message Notification User Guide</i> .

Parameter	Description
Auto Restart upon Exception	<p>Whether to enable automatic restart. If this function is enabled, any job that has become abnormal will be automatically restarted.</p> <p>If this option is selected, you need to set the following parameters:</p> <ul style="list-style-type: none">● Max. Retry Attempts: maximum number of retry times upon an exception. The unit is times/hour.<ul style="list-style-type: none">– Unlimited: The number of retries is unlimited.– Limited: The number of retries is user-defined.● Restore Job from Checkpoint: Restore the job from the saved checkpoint. <p>NOTE This parameter cannot be configured for Flink SQL jobs or Flink OpenSource SQL jobs.</p> <p>If this parameter is selected, you need to set Checkpoint Path for Flink Jar jobs.</p> <p>Checkpoint Path: Select the checkpoint saving path. The checkpoint path must be the same as that you set in the application package. Note that the checkpoint path for each job must be unique. Otherwise, the checkpoint cannot be obtained.</p>

3.3.8 Flink Job Details

After creating a job, you can view the job details to learn about the following information:

- [Viewing Job Details](#)
- [Checking the Job Monitoring Information](#)
- [Viewing the Task List of a Job](#)
- [Viewing the Job Execution Plan](#)
- [Viewing Job Submission Logs](#)
- [Viewing Job Running Logs](#)
- [Viewing Job Tags](#)

Viewing Job Details

This section describes how to view job details. After you create and save a job, you can click the job name to view job details, including SQL statements and parameter settings. For a Jar job, you can only view its parameter settings.

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Click the name of the job to be viewed. The **Job Detail** tab is displayed.

In the **Job Details** tab, you can view SQL statements, configured parameters, and total cost for the job.

The following uses a Flink SQL job as an example.

Table 3-13 Description

Parameter	Description
Type	Job type, for example, Flink SQL
Name	Flink job name
Description	Description of a Flink job
Status	Running status of a job
Running Mode	If your job runs on a shared queue, this parameter is Shared . If your job runs on a custom queue with dedicated resources, this parameter is Exclusive .
Flink Version	Version of Flink selected for the job.
Runtime Configuration	Displayed when a user-defined parameter is added to a job
CUs	Number of CUs configured for a job
Job Manager CUs	Number of job manager CUs configured for a job.
Parallelism	Number of jobs that can be concurrently executed by a Flink job
CU(s) per TM	Number of CUs occupied by each Task Manager configured for a job
Slot(s) per TM	Number of Task Manager slots configured for a job
OBS Bucket	OBS bucket name. After Enable Checkpointing and Save Job Log are enabled, checkpoints and job logs are saved in this bucket.
Save Job Log	Whether the job running logs are saved to OBS
Alarm Generation upon Job Exception	Whether job exceptions are reported
SMN Topic	Name of the SMN topic. This parameter is displayed when Alarm Generation upon Job Exception is enabled.
Auto Restart upon Exception	Whether automatic restart is enabled.
Max. Retry Attempts	Maximum number of retry times upon an exception. Unlimited means the number is not limited.

Parameter	Description
Savepoint	OBS path of the savepoint
Enable Checkpointing	Whether checkpointing is enabled
Checkpoint Interval	Interval between storing intermediate job running results to OBS. The unit is second.
Checkpoint Mode	Checkpoint mode. Available values are as follows: <ul style="list-style-type: none"> ● At least once: Events are processed at least once. ● Exactly once: Events are processed only once.
Idle State Retention Time	Defines for how long the state of a key is retained without being updated before it is removed in GroupBy or Window.
Dirty Data Policy	Policy for processing dirty data. The value is displayed only when there is a dirty data policy. Available values are as follows: <ul style="list-style-type: none"> Ignore Trigger a job exception Save
Dirty Data Dump Address	OBS path for storing dirty data when Dirty Data Policy is set to Save .
Created	Time when a job is created
Updated	Time when a job was last updated

----End

Checking the Job Monitoring Information

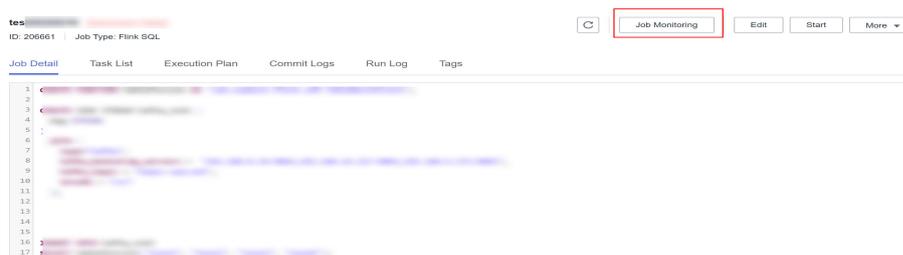
You can use Cloud Eye to view details about job data input and output.

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Click the name of the job you want. The job details are displayed.

Click **Job Monitoring** in the upper right corner of the page to switch to the Cloud Eye console.

Figure 3-12 Monitoring a Job



The following table describes monitoring metrics related to Flink jobs.

Table 3-14 Monitoring metrics related to Flink jobs

Name	Description
Flink Job Data Read Rate	Displays the data input rate of a Flink job for monitoring and debugging. Unit: record/s.
Flink Job Data Write Rate	Displays the data output rate of a Flink job for monitoring and debugging. Unit: record/s.
Flink Job Total Data Read	Displays the total number of data inputs of a Flink job for monitoring and debugging. Unit: records
Flink Job Total Data Write	Displays the total number of output data records of a Flink job for monitoring and debugging. Unit: records
Flink Job Byte Read Rate	Displays the number of input bytes per second of a Flink job. Unit: byte/s
Flink Job Byte Write Rate	Displays the number of output bytes per second of a Flink job. Unit: byte/s
Flink Job Total Read Byte	Displays the total number of input bytes of a Flink job. Unit: byte
Flink Job Total Write Byte	Displays the total number of output bytes of a Flink job. Unit: byte
Flink Job CPU Usage	Displays the CPU usage of Flink jobs. Unit: %
Flink Job Memory Usage	Displays the memory usage of Flink jobs. Unit: %
Flink Job Max Operator Latency	Displays the maximum operator delay of a Flink job. The unit is ms .
Flink Job Maximum Operator Backpressure	Displays the maximum operator backpressure value of a Flink job. A larger value indicates severer backpressure. 0 : OK 50 : low 100 : high

----End

Viewing the Task List of a Job

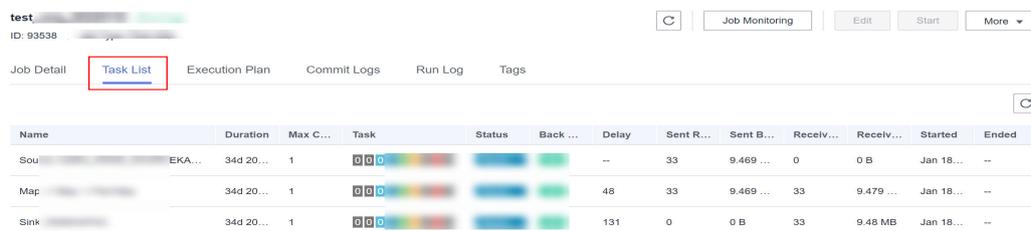
You can view details about each task running on a job, including the task start time, number of received and transmitted bytes, and running duration.

NOTE

If the value is **0**, no data is received from the data source.

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** Click the name of the job you want. The job details are displayed.
- Step 3** On **Task List** and view the node information about the task.

Figure 3-13 Task list



View the operator task list. The following table describes the task parameters.

Table 3-15 Parameter description

Parameter	Description
Name	Name of an operator.
Duration	Running duration of an operator.
Max Concurrent Jobs	Number of parallel tasks in an operator.
Task	Operator tasks are categorized as follows: <ul style="list-style-type: none"> • The digit in red indicates the number of failed tasks. • The digit in light gray indicates the number of canceled tasks. • The digit in yellow indicates the number of tasks that are being canceled. • The digit in green indicates the number of finished tasks. • The digit in blue indicates the number of running tasks. • The digit in sky blue indicates the number of tasks that are being deployed. • The digit in dark gray indicates the number of tasks in a queue.
Status	Status of an operator task.

Parameter	Description
Back Pressure Status	Working load status of an operator. Available options are as follows: <ul style="list-style-type: none"> ● OK: indicates that the operator is in normal working load. ● LOW: indicates that the operator is in slightly high working load. DLI processes data quickly. ● HIGH: indicates that the operator is in high working load. The data input speed at the source end is slow.
Delay	Duration from the time when source data starts being processed to the time when data reaches the current operator. The unit is millisecond.
Sent Records	Number of data records sent by an operator.
Sent Bytes	Number of bytes sent by an operator.
Received Bytes	Number of bytes received by an operator.
Received Records	Number of data records received by an operator.
Started	Time when an operator starts running.
Ended	Time when an operator stops running.

----End

Viewing the Job Execution Plan

You can view the execution plan to understand the operator stream information about the running job.

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** Click the name of the job you want. The job details are displayed.
- Step 3** Click the **Execution Plan** tab to view the operator flow direction.

Figure 3-14 Execution plan



Click a node. The corresponding information is displayed on the right of the page.

- Scroll the mouse wheel to zoom in or out.

- The stream diagram displays the operator stream information about the running job in real time.

----End

Viewing Job Submission Logs

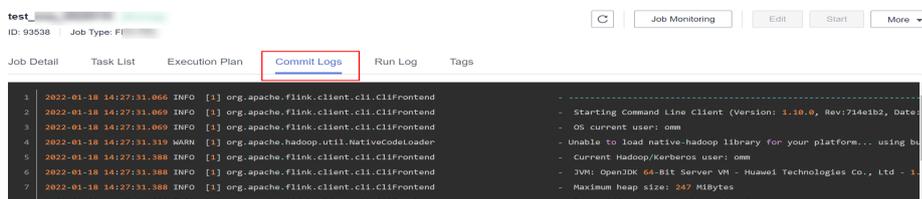
You can view the submission logs to locate the fault.

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Click the name of the job you want. The job details are displayed.

Step 3 In the **Commit Logs** tab, view the information about the job submission process.

Figure 3-15 Commit logs



----End

Viewing Job Running Logs

You can view the run logs to locate the faults occurring during job running.

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 Click the name of the job you want. The job details are displayed.

Step 3 On the **Run Log** tab page, you can view the **Job Manager** and **Task Manager** information of the running job.

Figure 3-16 Run logs



Information about JobManager and TaskManager is updated every minute. Only run logs of the last minute are displayed by default.

If you select an OBS bucket for saving job logs during the job configuration, you can switch to the OBS bucket and download log files to view more historical logs.

If the job is not running, information on the **Task Manager** page cannot be viewed.

----End

Viewing Job Tags

You can view, add, modify, and delete job tags.

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** Click the name of the job to be viewed. The **Job Details** page is displayed.
- Step 3** On the **Tags** tab page, view the tag information about the current job.



For more information about job tags, see [Tag Management](#).

----End

3.3.9 Tag Management

A tag is a key-value pair customized by users and used to identify cloud resources. It helps users to classify and search for cloud resources. A tag consists of a tag key and a tag value.

DLI allows you to add tags to Flink jobs. You can add tags to Flink jobs to identify information such as the project name, service type, and background. If you use tags in other cloud services, you are advised to create the same tag key-value pairs for cloud resources used by the same business to keep consistency.

DLI supports the following two types of tags:

- Resource tags: indicate non-global tags created on DLI.
- Predefined tags: global tags created on Tag Management Service (TMS).
For more information about predefined tags, see [Tag Management Service User Guide](#).

This section includes the following content:

- [Managing a Job Tag](#)
- [Searching for a Job by Tag](#)

Managing a Job Tag

DLI allows you to add, modify, or delete tags for jobs.

- Step 1** In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
- Step 2** Click the name of the job to be viewed. The **Job Details** page is displayed.
- Step 3** Click **Tags** to display the tag information about the current job.

Figure 3-17 Managing a job tag



Step 4 Click **Add/Edit Tag** to open to the **Add/Edit Tag** dialog box.

Step 5 Configure the tag parameters in the **Add/Edit Tag** dialog box.

Figure 3-18 Adding a tag

Table 3-16 Tag parameters

Parameter	Description
Tag key	<p>You can perform the following operations:</p> <ul style="list-style-type: none"> Click the text box and select a predefined tag key from the drop-down list. To add a predefined tag, you need to create one on TMS and then select it from the Tag key drop-down list. You can click View predefined tags to go to the Predefined Tags page of the TMS console. Then, click Create Tag in the upper corner of the page to create a predefined tag. For details, see Creating Predefined Tags in <i>Tag Management Service User Guide</i>. Enter a tag key in the text box. <p>NOTE A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed, but the value cannot start or end with a space or start with _sys_.</p>

Parameter	Description
Tag value	<p>You can perform the following operations:</p> <ul style="list-style-type: none"> Click the text box and select a predefined tag value from the drop-down list. Enter a tag value in the text box. <p>NOTE A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.::=+@) are allowed. The value cannot start or end with a space.</p>

NOTE

- A maximum of 20 tags can be added.
- Only one tag value can be added to a tag key.
- The key name in each resource must be unique.

Step 6 Click **OK**.

Step 7 (Optional) In the tag list, locate the row where the tag you want to delete resides, click **Delete** in the **Operation** column to delete the tag.

----End

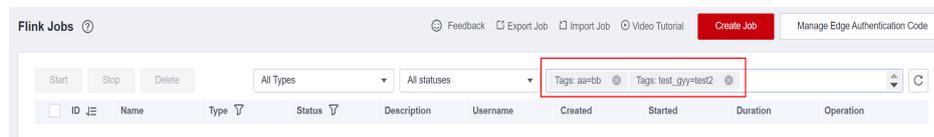
Searching for a Job by Tag

If tags have been added to a job, you can search for the job by setting tag filtering conditions to quickly find it.

Step 1 In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.

Step 2 In the upper right corner of the page, click the search box and select **Tags**.

Figure 3-19 Searching for jobs by tag



Step 3 Choose a tag key and value as prompted. If no tag key or value is available, create a tag for the job. For details, see [Managing a Job Tag](#).

Step 4 Choose other tags to generate a tag combination for job search. You can add a maximum of 10 tags to search for jobs. If you specify more than one tag, jobs containing all the specified tags will be displayed.

Step 5 Click search icon. The target job will be displayed in the job list.

----End

3.4 Spark Job Management

3.4.1 Spark Job Management

Based on the open-source Spark, DLI optimizes performance and reconstructs services to be compatible with the Apache Spark ecosystem and interfaces, and executes batch processing tasks.

DLI also allows you to use Spark jobs to access DLI metadata.

Spark job management provides the following functions:

- [Creating a Spark Job](#)
- [Re-executing a Job](#)
- [Searching for a Job](#)
- [Terminating a Job](#)

In addition, you can click **Quick Links** to switch to the details on *User Guide*.

Spark Jobs Page

On the **Overview** page, click **Spark Jobs** to go to the SQL job management page. Alternatively, you can click **Job Management > Spark Jobs**. The page displays all Spark jobs. If there are a large number of jobs, they will be displayed on multiple pages. DLI allows you to view jobs in all statuses.

Table 3-17 Job management parameters

Parameter	Description
Job ID	ID of a submitted Spark job, which is generated by the system by default.
Name	Name of a submitted Spark job.
Queues	Queue where the submitted Spark job runs
Username	Name of the user who executed the Spark job
Status	Job status. The following values are available: <ul style="list-style-type: none">• Starting: The job is being started.• Running: The job is being executed.• Failed: The session has exited.• Finished: The session is successfully executed.• Restoring: The job is being restored.
Created	Time when a job is created. Jobs can be displayed in ascending or descending order of the job creation time.
Last Modified	Time when a job is completed.

Parameter	Description
Operation	<ul style="list-style-type: none">● Edit: You can modify the current job configuration and re-execute the job.● SparkUI: After you click this button, the Spark job execution page is displayed. <p>NOTE</p> <ul style="list-style-type: none">● The SparkUI page cannot be viewed for jobs in the Starting state.● Currently, only the latest 100 job information records are displayed on the SparkUI of DLI. <ul style="list-style-type: none">● Terminate Job: Cancel a job that is being started or running.● Re-execute: Run the job again.● Archive Log: Save job logs to the temporary bucket created by DLI.● Commit Log: View the logs of submitted jobs.● Driver Log: View the logs of running jobs.

Re-executing a Job

On the **Spark Jobs** page, click **Edit** in the **Operation** column of the job. On the Spark job creation page that is displayed, modify parameters as required and execute the job.

Searching for a Job

On the **Spark Jobs** page, select **Status** or **Queues**. The system displays the jobs that meet the filter condition in the job list.

Terminating a Job

On the **Spark Jobs** page, choose **More > Terminate Job** in the **Operation** column of the job that you want to stop.

3.4.2 Creating a Spark Job

DLI provides fully-managed Spark computing services by allowing you to execute Spark jobs.

On the **Overview** page, click **Create Job** in the upper right corner of the **Spark Jobs** tab or click **Create Job** in the upper right corner of the **Spark Jobs** page. The Spark job editing page is displayed.

On the Spark job editing page, a message is displayed, indicating that a temporary DLI data bucket will be created. The created bucket is used to store temporary data generated by DLI, such as job logs and job results. You cannot view job logs if you choose not to create it. The bucket will be created and the default bucket name is used.

If you do not need to create a DLI temporary data bucket and do not want to receive this message, select **Do not show again** and click **Cancel**.

Prerequisites

- You have uploaded the dependencies to the corresponding OBS bucket on the **Data Management > Package Management** page. For details, see [Creating a Package](#).
- Before creating a Spark job to access other external data sources, such as OpenTSDB, HBase, Kafka, GaussDB(DWS), RDS, CSS, CloudTable, DCS Redis, and DDS MongoDB, you need to create a cross-source connection to enable the network between the job running queue and external data sources.
 - For details about the external data sources that can be accessed by Spark jobs, see [Cross-Source Analysis Development Methods](#).
 - For details about how to create a datasource connection, see [Enhanced Datasource Connections](#).

On the **Resources > Queue Management** page, locate the queue you have created, and choose **More > Test Address Connectivity** in the **Operation** column to check whether the network connection between the queue and the data source is normal. For details, see [Testing Address Connectivity](#).

Procedure

1. In the left navigation pane of the DLI management console, choose **Job Management > Spark Jobs**. The **Spark Jobs** page is displayed.
Click **Create Job** in the upper right corner. In the job editing window, you can set parameters in **Fill Form** mode or **Write API** mode.
The following uses the **Fill Form** as an example. In **Write API** mode, refer to the *Data Lake Insight API Reference* for parameter settings.
2. Select a queue.
Select the queue you want to use from the drop-down list box.
3. Select a Spark version. Select a supported Spark version from the drop-down list. The latest version is recommended.
4. Configure the job.
Configure job parameters by referring to [Table 3-18](#).

Figure 3-20 Spark job configuration

The screenshot shows a configuration form for a Spark job. The fields and their values are as follows:

- Job Name (--name):** test-catoonfile
- Application:** obs://ai-test/wwwTest/catConfigFile.py
- Main Class (--class):** This parameter cannot be left blank if the Application is a JAR package.
- Application Parameters:** Specify each parameter on a separate line.
- Spark Arguments (--conf):** Enter arguments using the key = value format. Press Enter to separate multiple key-value pairs.
- Job Type:** Basic (selected), AI-enhanced, Image
- JAR Package Dependencies (--jars):** Use the DLI program package for production services
- Python File Dependencies (--py-files):** Use the DLI program package for production services
- Other Dependencies (--files):** Use the DLI program package for production services
- Group Name:** Use the DLI program package for production services
- Access Metadata:** No (selected), Yes
- Retry upon Failure:** No (selected), Yes
- Advanced Settings:** Skip, Configure (selected)

Table 3-18 Job configuration parameters

Parameter	Description
Job Name (--name)	Set a job name.
Application	Select the package to be executed. The value can be .jar or .py .
Main Class (--class)	Enter the name of the main class. When the application type is .jar , the main class name cannot be empty.
Application Parameters	User-defined parameters. Separate multiple parameters by Enter . These parameters can be replaced with global variables. For example, if you create a global variable batch_num on the Global Configuration > Global Variables page, you can use {{batch_num}} to replace a parameter with this variable after the job is submitted.

Parameter	Description
Spark Arguments (--conf)	<p>Enter a parameter in the format of key=value. Press Enter to separate multiple key-value pairs. For details about the parameters, see Spark Configuration.</p> <p>These parameters can be replaced with global variables. For example, if you create a global variable custom_class on the Global Configuration > Global Variables page, you can use "spark.sql.catalog"={{custom_class}} to replace a parameter with this variable after the job is submitted.</p> <p>NOTE</p> <ul style="list-style-type: none"> The JVM garbage collection algorithm cannot be customized for Spark jobs. If the Spark version is 3.1.1, configure Spark parameters (--conf) to select a dependent module. For details about the example configuration, see Table 3-19.
Job Type	<p>Set this parameter when you select a CCE queue. Type of the Spark image used by a job. The options are as follows:</p> <ul style="list-style-type: none"> Basic: Basic images provided by DLI. Select this option for non-AI jobs. AI-enhanced: AI images provided by DLI. Select this option for AI jobs. Image: Custom Spark images. Select an existing image name and version on SWR.
JAR Package Dependencies (--jars)	<p>JAR file on which the Spark job depends. You can enter the JAR package name or the corresponding OBS path. The format is as follows: obs://Bucket name/Folder name/Package name.</p>
Python File Dependencies (--py-files)	<p>py-files on which the Spark job depends. You can enter the Python file name or the corresponding OBS path of the Python file. The format is as follows: obs://Bucket name/Folder name/File name.</p>
Other Dependencies (--files)	<p>Other files on which the Spark job depends. You can enter the name of the dependency file or the corresponding OBS path of the dependency file. The format is as follows: obs://Bucket name/Folder name/File name.</p>
Group Name	<p>If you select a group when creating a package, you can select all the packages and files in the group. For details about how to create a package, see Creating a Package.</p>
Access Metadata	<p>Whether to access metadata through Spark jobs. .</p>
Retry upon Failure	<p>Indicates whether to retry a failed job.</p> <p>If you select Yes, you need to set the following parameters:</p> <p>Maximum Retries: Maximum number of retry times. The maximum value is 100.</p>

Parameter	Description
Advanced Settings	<ul style="list-style-type: none"> • Skip • Configure <ul style="list-style-type: none"> – Select Dependency Resources: For details about the parameters, see Table 3-20. – Configure Resources: For details about the parameters, see Table 3-21.

Table 3-19 Spark Parameter (--conf) configuration

Datasource	Example Value
CSS	spark.driver.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/css/* spark.executor.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/css/*
DWS	spark.driver.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/dws/* spark.executor.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/dws/*
HBase	spark.driver.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/hbase/* spark.executor.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/hbase/*
OpenTSDB	spark.driver.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/opentsdb/* spark.executor.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/opentsdb/*
RDS	spark.driver.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/rds/* spark.executor.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/rds/*
Redis	spark.driver.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/redis/* spark.executor.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/redis/*
Mongo	spark.driver.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/mongo/* spark.executor.extraClassPath=/usr/share/extension/dli/spark-jar/datasource/mongo/*

Figure 3-21 Creating a Spark job - advanced settings

Select Dependency Resources

Module Name

Resource Package

Configure Resources

Resource Specifications

Start by selecting one of the resource sets, but then you can modify the executors, executor memory, the number of driver cores, and the driver memory as needed.

Executor Memory

Executor Cores

Executors

Driver Cores

Driver Memory

5. Set the following parameters in advanced settings:
 - **Select Dependency Resources:** For details about the parameters, see [Table 3-20](#).
 - **Configure Resources:** For details about the parameters, see [Table 3-21](#).

NOTE

The parallelism degree of Spark resources is jointly determined by the number of Executors and the number of Executor CPU cores.

Maximum number of tasks that can be concurrently executed = Number of Executors x Number of Executor CPU cores

You can properly plan compute resource specifications based on the compute CUs of the queue you have purchased.

Note that Spark tasks need to be jointly executed by multiple roles, such as **driver** and **executor**. So, the number of executors multiplied by the number of executor CPU cores must be less than the number of compute CUs of the queue to prevent other roles from failing to start Spark tasks. For more information about roles for Spark tasks, see [Apache Spark](#).

Calculation formula for Spark job parameters:

- CUs = Driver Cores + Executors x Executor Cores
- Memory = Driver Memory + (Executors x Executor Memory)

Table 3-20 Parameters for selecting dependency resources

Parameter	Description
modules	<p>If the Spark version is 3.1.1, you do not need to select a module. Configure Spark parameters (--conf).</p> <p>Dependency modules provided by DLI for executing datasource connection jobs. To access different services, you need to select different modules.</p> <ul style="list-style-type: none"> • CloudTable/MRS HBase: sys.datasource.hbase • DDS: sys.datasource.mongo • CloudTable/MRS OpenTSDB: sys.datasource.opentsdb • DWS: sys.datasource.dws • RDS MySQL: sys.datasource.rds • RDS PostGre: sys.datasource.rds • DCS: sys.datasource.redis • CSS: sys.datasource.css <p>DLI internal modules include:</p> <ul style="list-style-type: none"> • sys.res.dli-v2 • sys.res.dli • sys.datasource.dli-inner-table
Resource Package	JAR package on which the Spark job depends.

Table 3-21 Resource specification parameters

Parameter	Description
Resource Specifications	<p>Select a resource specification from the drop-down list box. The system provides three resource specification options for you to choose from.</p> <p>Resource specifications involve the following parameters:</p> <ul style="list-style-type: none"> • Executor Memory • Executor Cores • Executors • Driver Cores • Driver Memory <p>If modified, your modified settings of the items are used.</p>

Parameter	Description
Executor Memory	Customize the configuration item based on the selected resource specifications. Memory of each Executor. It is recommended that the ratio of Executor CPU cores to Executor memory be 1:4.
Executor Cores	Number of CPU cores of each Executor applied for by Spark jobs, which determines the capability of each Executor to execute tasks concurrently.
Executors	Number of Executors applied for by a Spark job
Driver Cores	Number of CPU cores of the driver
Driver Memory	Driver memory size. It is recommended that the ratio of the number of driver CPU cores to the driver memory be 1:4.

6. Click **Execute** in the upper right corner of the Spark job editing page.
After the message "Batch processing job submitted successfully" is displayed, you can view the status and logs of the submitted job on the **Spark Jobs** page.

4 Queue Management

4.1 Overview

Queue

Queues in DLI are computing resources, which are the basis for using DLI. All executed jobs require computing resources.

Currently, DLI provides two types of queues: **For SQL** and **For general purpose**.

- For SQL: The queue is used to run SQL jobs.
- For general purpose: The queue is used to run Spark programs, Flink SQL jobs, and Flink Jar jobs.

Constraints

- A queue named **default** is preset in DLI for you to experience. Resources are allocated on demand. You are billed based on the amount of data scanned in each job (unit: GB).
- Queue types:
 - For SQL: Spark SQL jobs can be submitted to SQL queues.
 - For general purpose: The queue is used to run Spark programs, Flink SQL jobs, and Flink Jar jobs.

The queue type cannot be changed. If you want to use another queue type, purchase a new queue.

- The billing mode of a queue cannot be changed.
- The region of a queue cannot be changed.
- Queues with 16 CUs do not support scale-out or scale-in.
- Queues with 64 CUs do not support scale-in.
- When creating a queue, you can only select cross-AZ active-active for yearly/monthly queues and pay-per-use dedicated queues. The price of a cross-AZ queue is twice that of a single-AZ queue.
- A newly created queue can be scaled in or out only after a job is executed on the queue.

- DLI queues cannot access the Internet.

For details about how to access the Internet from an elastic resource pool, see [Configuring the Connection Between a DLI Queue and a Data Source on the Internet](#).

Difference Between Computing and Storage Resources

Table 4-1 Difference between computing and storage resources

Resource	How to Obtain	Billing Mode	Function
Compute resource	Create queue on the DLI management console.	Billed based on the amount of data scanned or CUHs.	Used for executing queries.
Storage resource	DLI has a 5 GB quota.	Billed based on the amount of stored data.	Used for storing data in the database and DLI tables.

NOTE

- Storage resources are internal storage resources of DLI for storing database and DLI tables and represent the amount of data stored in DLI.
- By default, DLI provides a 5 GB quota for storage resources. If you need over 5 GB storage resources, submit a service ticket to apply for more storage resources.
- A queue named **default** is preset in DLI. DLI adopts the pay-per-use billing mode. If you are uncertain about the required queue capacity or have no available queue capacity to run queries, you can execute jobs using this queue.
- The **default** queue is used only for user experience. It may be occupied by multiple users at a time. Therefore, it is possible that you fail to obtain the resource for related operations. You are advised to use a self-built queue to execute jobs.

Dedicated Queue

Resources of a dedicated queue are not released when the queue is idle. That is, resources are reserved regardless of whether the queue is used. Dedicated queues ensure that resources exist when jobs are submitted. When purchasing a pay-per-use queue, you can select a dedicated queue. Dedicated queues are billed for 24 hours regardless of whether they are used.

Cross-AZ Queues

An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proofing, and electricity facilities. Within an

AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to allow you to build cross-AZ high-availability systems.

DLI dual-AZ queues improve data availability by creating a duplicate queue in the second AZ. You can continuously use DLI when one AZ is unavailable. The dual-AZ mode suits those who require high availability.

A DLI dual-AZ queue is created with the same compute resources in two AZs. For example, if you require 1,400 CUs, you can select 1,400 CUs and select the dual-AZ option when creating a queue. Then, DLI creates 1,400 CUs dedicated compute resources in two AZs. When one AZ is unavailable, the other AZ can properly process your compute tasks.

NOTE

- Currently, dual-AZ mode is supported only for pay-per-use dedicated queues. The default queue and the normal pay-per-use queues are not supported.

Elastic Scaling (Pay-per-Use Queues)

DLI allows you to flexibly scale in or out pay-per-use queues on demand. After a pay-per-use queue with specified specifications is created, you can scale it in and out as required.

To change the queue specifications, see [Elastic Scaling](#). Currently, only pay-per-use queues can be scaled. After scaling, you are still charged based on the usage duration of CUs.

NOTE

Scaling can be performed for a newly created queue only when jobs are running on this queue.

Scheduled Elastic Scaling

DLI allows you to schedule tasks for periodic queue scaling. After creating a queue, the scheduled scaling tasks can be executed.

You can schedule an automatic scale-out/scale-in based on service requirements. The system periodically triggers queue scaling. For details, see [Scheduling CU Changes](#).

- After scaling, you are still charged based on the usage duration of CUs. Currently, scheduled scaling tasks are available only for a queue with more than 64 CUs. That is, the minimum number of CUs in a queue is 64.
- Yearly/monthly queues can be scaled out through scheduled tasks only. The fees cover the prepaid yearly/monthly package and the expanded pay-per-use CUs. In other words, you will be billed on a pay-per-use basis for the resources not included in the yearly/monthly package. Currently, scheduled scaling tasks are available only for a yearly/monthly queue with more than 64 CUs.

NOTE

Scaling can be performed for a newly created queue only when jobs are running on this queue.

Automatic Queue Scaling

Flink jobs use pay-per-use queues. DLI can automatically trigger scaling for jobs based on the job size.

NOTE

Scaling can be performed for a newly created queue only when there are jobs running on this queue.

Queue Management Page

Queue Management provides the following functions:

- [Managing Permissions](#)
- [Creating a Queue](#)
- [Deleting a Queue](#)
- [Modifying CIDR Block](#)
- [Elastic Scaling](#)
- [Scheduling CU Changes](#)
- [Testing Address Connectivity](#)
- [Creating a Topic for Key Event Notifications](#)
- [Managing Queue Tags](#)

NOTE

To receive notifications when a DLI job fails, **SMN Administrator** permissions are required.

The queue list displays all queues created by you and the **default** queue. You can view information about queues, such as the queue capacity, billing mode, and more. Queues are listed in chronological order by default in the queue list, with the most recently created queues displayed at the top.

Table 4-2 Parameter description

Parameter	Description
Name	Name of a queue.
Type	Queue type. <ul style="list-style-type: none">• For SQL• For general purpose• Spark queue (compatible with earlier versions)

Parameter	Description
Specifications	<p>Queue size. Unit: CU</p> <ul style="list-style-type: none"> • Pay-per-use queue: number of CUs based on which you are billed pay-per-use • Yearly/Monthly queue: number of CUs based on which you are billed yearly/monthly <p>CU is the pricing unit of queues. A CU consists of 1 vCPU and 4-GB memory. The computing capabilities of queues vary with queue specifications. The higher the specifications, the stronger the computing capability.</p>
Actual CUs	Actual size of the current queue.
Elastic Scaling	Target CU value for scheduled scaling, or the maximum and minimum CU values of the current specifications.
Billing Mode	<ul style="list-style-type: none"> • You are billed for the executed SQL jobs based on either of the following methods: <ul style="list-style-type: none"> – Pay-per-use/By CUH: The compute usage is billed based on the CUH used. It is recommended that you select CUH as the package type. – Pay-per-use/By SQL computations: You are billed based on the amount of data scanned of each job. It is recommended that you select Scanned data volume as the package type. <p>NOTE Only the default queue is billed based on the amount of data scanned. All user-defined queues are billed based on the CUH used.</p> <ul style="list-style-type: none"> – In dedicated resource mode of pay-per-use queues, you will be billed by calendar hour since the queue is created. In this mode, you can create enhanced datasource connections.
Username	Queue owner
Enterprise Project	<p>Displays the enterprise project to which the created queue belongs. If the project is not an enterprise project, -- is displayed.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. For details about how to set enterprise projects, see Enterprise Center Overview.</p> <p>NOTE This parameter is displayed only for users who have enabled the Enterprise Management Service.</p>
Description	Description of a queue specified during queue creation. If no description is provided, -- is displayed.

Parameter	Description
Operation	<ul style="list-style-type: none">● Delete: Allow you to delete the selected queue. You cannot delete a queue where there are running jobs or jobs are being submitted. NOTE Only pay-per-use queues can be deleted.● Manage Permissions: You can view the user permissions corresponding to the queue and grant permissions to other users.● More<ul style="list-style-type: none">– Restart: Forcibly restart a queue. NOTE Only the SQL queue has the Restart operation.– Modify Enterprise Project: Modify the enterprise project to which the selected queue belongs. NOTE This parameter is displayed only for users who have enabled the Enterprise Management Service. For details about how to set enterprise projects, see Enterprise Center Overview.– Elastic Scaling: You can select Scale-out or Scale-in as required. The number of CUs after modification must be an integer multiple of 16.– Modify Queue Specifications: You can select Scale-out or Scale-in as required. The number of CUs after modification must be an integer multiple of 16.– Schedule CU Changes: You can set different queue sizes at different time or in different periods based on the service period or usage. The system automatically performs scale-out or scale-in as scheduled. The After Modification value must be an integer multiple of 16.– Modifying CIDR Block: When DLI enhanced datasource connection is used, the CIDR block of the DLI queue cannot overlap with that of the data source. You can modify the CIDR block as required. Recommended CIDR block: 10.0.0.0~10.255.0.0/8~24 172.16.0.0~172.31.0.0/12~24– Test Address Connectivity: Test whether the queue is reachable to the specified address. The domain name and IP address are supported. The port can be specified.– Tag: You can add, edit, and delete tags.

4.2 Queue Permission Management

Scenario

- You can isolate queues allocated to different users by setting permissions to ensure data query performance.
- The administrator and queue owner have all permissions, which cannot be set or modified by other users.
- When setting queue permissions for a new user, ensure that the region of the user group to which the user belongs has the **Tenant Guest** permission. For details about the Tenant Guest permission and how to apply for the permission, see [Creating a User Group and Assigning Permissions](#) and [System Permissions](#).

Operations

Step 1 On the top menu bar of the DLI management console, click **Resources > Queue Management**.

Step 2 Select the queue to be configured and choose **Manage Permissions** in the **Operation** column. The **User Permission Info** area displays the list of users who have permissions on the queue.

You can assign queue permissions to new users, modify permissions for users who have some permissions of a queue, and revoke all permissions of a user on a queue.

- Assign permissions to a new user.
A new user does not have permissions on the queue.
 - a. Click **Set Permission** on the right of **User Permissions** page. The **Set Permission** dialog box is displayed.
 - b. Specify **Username** and select corresponding permissions.
 - c. Click **OK**.

[Table 4-3](#) describes the related parameters.

Figure 4-1 Queue permission granting

Grant Permission

* Username

Select the permissions to be granted to the user

<input type="checkbox"/> Select all	<input type="checkbox"/> Submit Jobs	<input type="checkbox"/> Terminate Job
<input type="checkbox"/> Delete Queues	<input type="checkbox"/> Revoke Permission	<input type="checkbox"/> View Other User's Permissions
<input type="checkbox"/> Grant Permission	<input type="checkbox"/> Modify Queue Specifications	
<input type="checkbox"/> Restart Queues		

Table 4-3 Parameter description

Parameter	Description
Username	Name of the authorized user. NOTE The username is an existing IAM user name and has logged in to the DLI management console.
Permission Settings	<ul style="list-style-type: none">▪ Delete Queues: This permission allows you to delete the queue.▪ Submit Jobs: This permission allows you to submit jobs using this queue.▪ Terminate Jobs: This permission allows you to terminate jobs submitted using this queue.▪ Grant Permission: This permission allows you to grant queue permissions to other users.▪ Revoke Permission: This permission allows you to revoke the queue permissions that other users have but cannot revoke the queue owner's permissions.▪ View Other User's Permissions: This permission allows you to view the queue permissions of other users.▪ Restart Queues: This permission allows you to restart queues.▪ Modify Queue Specifications: This permission allows you to modify queue specifications.

- To assign or revoke permissions of a user who has some permissions on the queue, perform the following steps:
 - a. In the list under **User Permission Info** for a queue, select the user whose permissions need to be modified and click **Set Permission** in the **Operation** column.
 - b. In the displayed **Set Permission** dialog box, modify the permissions of the current user. [Table 4-3](#) lists the detailed permission descriptions.
If all options under **Set Permission** are gray, you are not allowed to change permissions on this queue. You can apply to the administrator, queue owner, or other authorized users for queue permission granting and revoking.

Figure 4-2 Setting queue permissions

Set Permission

* Username

Select the permissions to be granted to the user

Select all

Delete Queues Submit Jobs Terminate Job

Grant Permission Revoke Permission View Other User's Permissions

Restart Queues Modify Queue Specifications

- c. Click **OK**.
- To revoke all permission of a user on the queue, perform the following steps:
In the user list under **Permission Info**, select the user whose permission needs to be revoked and click **Revoke Permission** under **Operation**. In the **Revoke Permission** dialog box, click **OK**. All permissions on this queue are revoked.

----End

4.3 Creating a Queue

Before executing a job, you need to create a queue.

NOTE

- If you use a sub-account to create a queue for the first time, log in to the DLI management console using the main account and keep records in the DLI database before creating a queue.
- It takes 6 to 10 minutes for a job running on a new queue for the first time.
- After a queue is created, if no job is run within one hour, the system releases the queue.
- Queues with 16 CUs do not support scale-out or scale-in.
- Queues with 64 CUs do not support scale-in.

Procedure

1. You can create a queue on the **Overview**, **SQL Editor**, or **Queue Management** page.
 - In the upper right corner of the **Overview** page, click **Purchase Queue**.
 - To create a queue on the **Queue Management** page:
 - i. In the navigation pane of the DLI management console, choose **Resources > Queue Management**.
 - ii. In the upper right corner of the **Queue Management** page, click **Buy Queue** to create a queue.
 - To create a queue on the **SQL Editor** page:
 - i. In the navigation pane of the DLI management console, click **SQL Editor**.
 - ii. Click **Queues**. On the tab page displayed, click  on the right to create a queue.
2. On the **Buy Queue** page displayed, set the parameters according to [Table 4-4](#).

Table 4-4 Parameters

Parameter	Description
Billing Mode	Pay-per-use Billing for CUH used = Number of CUs x Usage duration x Unit price. You are billed for used CUs on an hourly basis (rounded up to the nearest hour). The pay-per-use billing mode is adopted. You are advised to purchase the CUH package to enjoy preferential price.
Region	Select a region. Select a region near you to ensure the lowest latency possible.
Project	Each region corresponds to a project.
Name	Name of a queue. <ul style="list-style-type: none">• The queue name can contain only digits, letters, and underscores (_), but cannot contain only digits, start with an underscore (_), or be left unspecified.• The length of the name cannot exceed 128 characters. NOTE The queue name is case-insensitive. Uppercase letters will be automatically converted to lowercase letters.
Type	<ul style="list-style-type: none">• For SQL: compute resources used for SQL jobs.• For general purpose: compute resources used for Spark and Flink jobs. NOTE Selecting Dedicated Resource Mode when Pay-per-use is selected for Billing Mode enables you to create a dedicated queue. Enhanced datasource connections can only be created for dedicated queues. You will be billed by calendar hour since the dedicated queue is created.
Specifications	The compute nodes' total number of CUs. One CU equals one vCPU and 4 GB of memory. DLI automatically assigns CPU and memory resources to each compute node, and the client does not need to know how many compute nodes are being used.
Enterprise Project	If the created queue belongs to an enterprise project, you can select the corresponding enterprise project. An enterprise project facilitates project-level management and grouping of cloud resources and users. For details about how to set an enterprise project, see Enterprise Management User Guide . NOTE This parameter is displayed only for users who have enabled the Enterprise Management Service.
Description	Description of the queue to be created. The description can contain a maximum of 128 characters.

Parameter	Description
Advanced Settings	<p>In the Queue Type area, select Dedicated Resource Mode and then click Advanced Settings.</p> <ul style="list-style-type: none">• Default: The system automatically configures the parameter.• Custom CIDR Block: You can specify the CIDR block. For details, see Modifying the CIDR Block. If DLI enhanced datasource connection is used, the CIDR block of the DLI queue cannot overlap with that of the data source. Recommended CIDR block: 10.0.0.0~10.255.0.0/8~24 172.16.0.0~172.31.0.0/12~24 192.168.0.0~192.168.0.0/16~24 Queue Type: When running an AI-related SQL job, select AI-enhanced. When running other jobs, select Basic.
Tags	<p>Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS). For details, see Tag Management Service User Guide.</p> <p>NOTE</p> <ul style="list-style-type: none">• A maximum of 20 tags can be added.• Only one tag value can be added to a tag key.• The key name in each resource must be unique.• Tag key: Enter a tag key name in the text box. <p>NOTE</p> <p>A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.::=+@) are allowed, but the value cannot start or end with a space or start with _sys_.</p> <ul style="list-style-type: none">• Tag value: Enter a tag value in the text box. <p>NOTE</p> <p>A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.::=+@) are allowed. The value cannot start or end with a space.</p>

3. Click **Buy Now** to confirm the configuration.
4. Confirm the configuration and click **Submit**.

If the queue name already exists, the system displays a message indicating that the queue name already exists when you click **Submit**. In this case, click **Previous** to go back to modify the queue name.

5. After a queue is created, you can view and select the queue for use on the **Queue Management** page.

 NOTE

It takes 6 to 10 minutes for a job running on a new queue for the first time.

4.4 Deleting a Queue

You can delete a queue based on actual conditions.

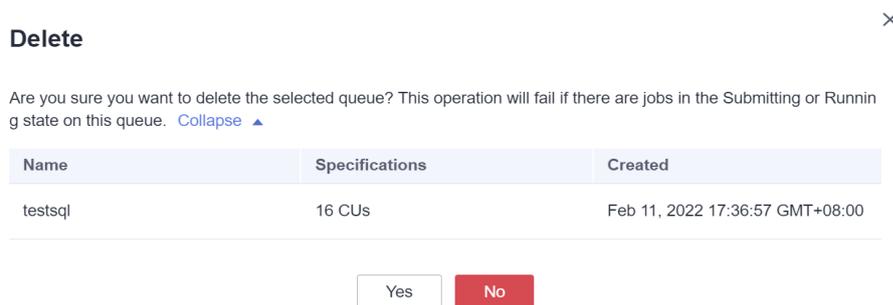
 NOTE

- This operation will fail if there are jobs in the **Submitting** or **Running** state on this queue.
- Deleting a queue does not cause table data loss in your database.

Procedure

- Step 1** On the left of the DLI management console, click **Resources >Queue Management**.
- Step 2** Locate the row where the target queue locates and click **Delete** in the **Operation** column.

Figure 4-3 Deleting a Queue

 NOTE

If **Delete** in the **Operation** column is gray, the current user does not have the permission of deleting the queue. You can apply to the administrator for the permission.

- Step 3** In the displayed dialog box, click **OK**.

----End

4.5 Allocating a Queue to an Enterprise Project

You can create enterprise projects matching the organizational structure of your enterprises to centrally manage cloud resources across regions by project. Then you can create user groups and users with different permissions and add them to enterprise projects.

DLI allows you to select an enterprise project when creating a queue. This section describes how to bind a DLI queue to and modify an enterprise project.

 **NOTE**

Currently, enterprise projects can be modified only for queues that have not been added to elastic resource pools.

Prerequisites

You have logged in to the Enterprise Project Management Service console and created an enterprise project by referring to [Creating an Enterprise Project](#).

Binding an Enterprise Project

When creating a DLI queue, you can select a created enterprise project for **Enterprise Project**.

Alternatively, you can click **Create Enterprise Project** to go to the Enterprise Project Management Service console to create an enterprise project and view existing ones.

For details about how to create a queue, see [Creating a Queue](#).

Modifying an Enterprise Project

You can modify the enterprise project bound to a created cluster as needed.

1. Log in to the DLI management console.
2. In the navigation pane on the left, choose **Resources > Queue Management**.
3. In the queue list, locate the queue for which you want to modify the enterprise project, click **More** in the **Operation** column, and select **Modify Enterprise Project**.
4. In the **Modify Enterprise Project** dialog box displayed, select an enterprise project.
Alternatively, you can click **Create Enterprise Project** to go to the Enterprise Project Management Service console to create an enterprise project and view existing ones.
5. After the modification, click **OK** to save the enterprise project information of the queue.

4.6 Modifying the CIDR Block

If the CIDR block of the DLI queue conflicts with that of the user data source, you can change the CIDR block of the queue.

 **NOTE**

If the queue whose CIDR block is to be modified has jobs that are being submitted or running, or the queue has been bound to enhanced datasource connections, the CIDR block cannot be modified.

Procedure

NOTE

Currently, you can modify the CIDR block of for queues whose billing mode is **Yearly/Monthly** or **Pay-per-use/Dedicated resource mode**.

1. On the left of the DLI management console, click **Resources >Queue Management**.
2. Select the queue to be modified and click **Modify CIDR Block** in the **Operation** column.

Figure 4-4 Modifying a CIDR block

×

Modify CIDR Block

Name testaaa

CIDR Block / ▼

If you need to use DLI enhanced datasource connections, the CIDR block entered here cannot be the same as that of the data source.
Recommended CIDR blocks:
10.0.0.0~10.255.0.0/8~24, 172.16.0.0~172.31.0.0/12~24, 192.168.0.0~192.168.0.0/16~24

3. Enter the required CIDR block and click **OK**. After the CIDR block of the queue is successfully changed, wait for 5 to 10 minutes until the cluster to which the queue belongs is restarted and then run jobs on the queue.

Recommended CIDR block:

10.0.0.0~10.255.0.0/8~24

172.16.0.0~172.31.0.0/12~24

192.168.0.0~192.168.0.0/16~24

4.7 Elastic Queue Scaling

Prerequisites

Elastic scaling can be performed for a newly created pay-per-use queue only when there were jobs running in this queue.

Precautions

- Queues with 16 CUs do not support scale-out or scale-in.
- Queues with 64 CUs do not support scale-in.
- Only queues whose billing mode is **Pay-per-use/By CUH** and **Pay-per-use/Dedicated resource mode** support elastic scaling.

- If **Status of queue xxx is assigning, which is not available** is displayed on the **Elastic Scaling** page, the queue can be scaled only after the queue resources are allocated.

Scaling Out

If the current queue specifications do not meet service requirements, you can add the number of CUs to scale out the queue.

NOTE

Scale-out is time-consuming. After you perform scale-out on the **Elastic Scaling** page of DLI, wait for about 10 minutes. The duration is related to the CU amount to add. After a period of time, refresh the **Queue Management** page and check whether values of **Specifications** and **Actual CUs** are the same to determine whether the scale-out is successful. Alternatively, on the **Job Management** page, check the status of the **SCALE_QUEUE** SQL job. If the job status is **Scaling**, the queue is being scaled out.

The procedure is as follows:

1. On the left of the DLI management console, click **Resources > Queue Management**.
2. Select the queue to be scaled out and choose **More > Elastic Scaling** in the **Operation** column.
3. On the displayed page, select **Scale-out** for **Operation** and set the scale-out amount.

Figure 4-5 Elastic scale-out

Name	sparktest
Billing Mode	Pay-per-use
Actual CUs	16 CUs
Opération	<input checked="" type="button" value="Scale-out"/> <input type="button" value="Scale-in"/>
Amount (CUs)	<input type="button" value="−"/> 16 <input type="button" value="+"/>
Final CU Count	32 CUs

4. Click OK.

Scaling In

If the current queue specifications are too much for your computing service, you can reduce the number of CUs to scale in the queue.

 NOTE

- Scale-in is time-consuming. After you perform scale-in on the **Elastic Scaling** page of DLI, wait for about 10 minutes. The duration is related to the CU amount to reduce. After a period of time, refresh the **Queue Management** page and check whether values of **Specifications** and **Actual CUs** are the same to determine whether the scale-in is successful. Alternatively, on the **Job Management** page, check the status of the **SCALE_QUEUE** SQL job. If the job status is **Scaling**, the queue is being scaled in.
- The system may not fully scale in the queue to the target size. If the current queue is in use or the service volume of the queue is large, the scale-in may fail or only partial specifications may be reduced.
- By default, the minimum number of CUs is **16**. That is, when the queue specifications are **16 CUs**, you cannot scale in the queue.

The procedure is as follows:

1. On the left of the DLI management console, click **Resources > Queue Management**.
2. Select the queue to be scaled out, click **More** in the **Operation** column, and select **Elastic Scaling**.
3. On the displayed page, select **Scale-in** for **Operation** and set the scale-in amount.

Figure 4-6 Manual scale-in

Name	createoutcuqueue
Billing Mode	Pay-per-use
Actual CUs	128 CUs
Opération	Scale-out Scale-in
Amount (CUs)	<input type="button" value="−"/> <input type="text" value="16"/> <input type="button" value="+"/>
Final CU Count	112 CUs

4. Click OK.

4.8 Scheduling CU Changes

Scenario

When services are busy, you might need to use more compute resources to process services in a period. After this period, you do not require the same amount of resources. If the purchased queue specifications are small, resources may be insufficient during peak hours. If the queue specifications are large, resources may be wasted.

DLI provides scheduled tasks for elastic scale-in and -out in the preceding scenario. You can set different queue sizes (CUs) at different time or in different periods

based on your service period or usage and the existing queue specifications to meet your service requirements and reduce costs.

Precautions

- Periodic scaling can be performed for a newly created queue only when there were jobs running in this queue.
- Scheduled scaling tasks are available only for a queue with more than 64 CUs. That is, the minimum specifications of a queue are 64 CUs.
- A maximum of 12 scheduled tasks can be created for each queue.
- When each scheduled task starts, the actual start time of the specification change has a deviation of 5 minutes. It is recommended that the task start time be at least 20 minutes earlier than the time when the queue is actually used.
- The interval between two scheduled tasks must be at least 2 hours.
- Changing the specifications of a queue is time-consuming. The time required for changing the specifications depends on the difference between the target specifications and the current specifications. You can view the specifications of the current queue on the **Queue Management** page.
- If a job is running in the current queue, the queue may fail to be scaled in to the target CU amount value. Instead, it will be scaled in to a value between the current queue specifications and the target specifications. The system will try to scale in again 1 hour later until the next scheduled task starts.
- If a scheduled task does not scale out or scale in to the target CU amount value, the system triggers the scaling plan again 15 minutes later until the next scheduled task starts.

Creating Periodic Task

- If only scale-out or scale-in is required, you need to create only one task for changing specifications. Set the **Task Name**, **Final CU Count**, and **Executed** parameters. For details, see [Table 4-5](#).
- To set both scale-out and scale-in parameters, you need to create two periodic tasks, and set the **Task Name**, **Final CU Count**, and **Executed** parameters. For details, see [Table 4-5](#).

The procedure is as follows:

1. On the left of the DLI management console, click **Resources > Queue Management**.
2. Locate the queue for which you want to schedule a periodic task for elastic scaling, and choose **More > Schedule CU Changes** in the **Operation** column.
3. On the displayed page, click **Create Periodic Task** in the upper right corner.
4. On the **Create Periodic Task** page, set the required parameters. Click **OK**.

Figure 4-7 Creating a periodic task

Create Periodic Task

* Task Name

Enable Task

Validity Period -

Total Number 32 CUs

* Final CU Count CUs

Repeat Select all
 Monday Tuesday Wednesday Thursday Friday Saturday
 Sunday

* Executed

Table 4-5 Parameter description

Parameter	Description
Task Name	Enter the name of the periodic task. <ul style="list-style-type: none"> The task name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_) or be left unspecified. The name can contain a maximum of 128 characters.
Enable Task	Whether to enable periodic elastic scaling. The task is enabled by default. If disabled, the task will not be triggered on time.
Validity Period	Time segment for executing the periodic task. The options include Date and Time .
Actual CUs	Queue specifications before scale-in or scale-out.
Final CUs	Specifications after the queue is scaled in or out. <p>NOTE</p> <ul style="list-style-type: none"> By default, the maximum specifications of a queue are 512 CUs. The minimum queue specifications for scheduled scaling are 64 CUs. That is, only when Actual CUs are more than 64 CUs, the scheduled scaling can be performed. The value of Actual CUs must be a multiple of 16.

Parameter	Description
Repeat	<p>Time when scheduled scale-out or scale-in is repeat. Scheduled tasks can be scheduled by week in Repeat.</p> <ul style="list-style-type: none">• By default, this parameter is not configured, indicating that the task is executed only once at the time specified by Executed.• If you select all, the plan is executed every day.• If you select some options of Repeat, the plan is executed once a week at all specified days. <p>NOTE</p> <ul style="list-style-type: none">• You do not need to set this parameter if you only need to perform scale-in or scale-out once.• If you have set scaling, you can set Repeat as required. You can also set the repeat period together with the validity period.
Executed	<p>Time when scheduled scale-out or scale-in is performed</p> <ul style="list-style-type: none">• When each scheduled task starts, the actual start time of the specification change has a deviation of 5 minutes. It is recommended that the task start time be at least 20 minutes earlier than the time when the queue is actually used.• The interval between two scheduled tasks must be at least 2 hours.

After a periodic task is created, you can view the specification change of the current queue and the latest execution time on the page for scheduling CU changes.

Alternatively, on the **Queue Management** page, check whether the **Specifications** change to determine whether the scaling is successful.

You can also go to the **Job Management** page and check the status of the **SCALE_QUEUE** job. If the job status is **Scaling**, the queue is being scaled in or out.

Modifying a Scheduled Task

If a periodic task cannot meet service requirements anymore, you can modify it on the **Schedule CU Changes** page.

1. In the navigation pane of the DLI management console, choose **Resources > Queue Management**.
2. Locate the queue for which you want to schedule a periodic task for elastic scaling, and choose **More > Schedule CU Changes** in the **Operation** column.
3. On the displayed page, click **Modify** in the **Operation** column. In the displayed dialog box, modify the task parameters as needed.

Deleting a Scheduled Task

If you do not need the task anymore, delete the task on the **Schedule CU Changes** page.

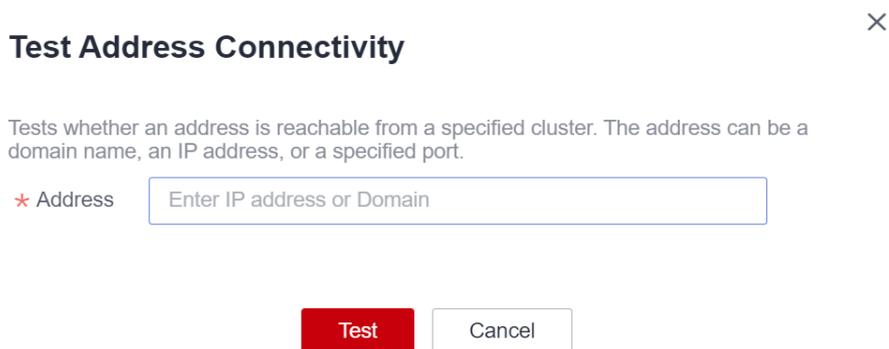
1. In the navigation pane of the DLI management console, choose **Resources >Queue Management**.
2. Locate the queue for which you want to schedule a periodic task for elastic scaling, and choose **More > Schedule CU Changes** in the **Operation** column.
3. On the displayed page, click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

4.9 Testing Address Connectivity

It can be used to test the connectivity between the DLI queue and the peer IP address specified by the user in common scenarios, or the connectivity between the DLI queue and the peer IP address bound to the datasource connection in cross-source connection scenarios. The operation is as follows:

1. On the **Queue Management** page, locate the row containing the target queue, click **More** in the **Operation** column, and select **Test Address Connectivity**.
2. On the **Test Address Connectivity** page, enter the address to be tested. The domain name and IP address are supported, and the port number can be specified.

Figure 4-8 Testing address connectivity



Test Address Connectivity ×

Tests whether an address is reachable from a specified cluster. The address can be a domain name, an IP address, or a specified port.

* Address

Test

3. Click **Test**.
If the test address is reachable, a message is displayed on the page, indicating that the address is reachable.
If the test address is unreachable, the system displays a message indicating that the address is unreachable. Check the network configurations and try again. Network configurations include the VPC peering and the datasource connection. Check whether they have been activated.

4.10 Creating an SMN Topic

Scenario

Once you have created an SMN topic, you can easily subscribe to it by going to the **Topic Management > Topics** page of the SMN console. You can choose to

receive notifications via SMS or email. After the subscription is successful, if a job fails, the system automatically sends a message to the subscription endpoint you specified.

- If a job fails within 1 minute of submission, a message notification is not triggered.
- If a job fails after 1 minute of submission, the system automatically sends a message to the subscriber terminal you specified.

Procedure

1. On the **Resources > Queue Management** page, click **Create SMN Topic** on the upper left side. The **Create SMN Topic** dialog box is displayed.

Figure 4-9 Creating an SMN topic

Create SMN Topic

After creating a message notification topic, you can subscribe to it in different ways. After the subscription succeeds, any job failure will automatically be sent to your subscription endpoints.

--Select-- ▼

2. Select a queue and click **OK**.

NOTE

- You can select a single queue or all queues.
- If you create a topic for a queue and another topic for all queues, the SMN of all queues does not include the message of the single queue.
- After a message notification topic is created, you will receive a message notification only when a Spark job created on the subscription queue fails.

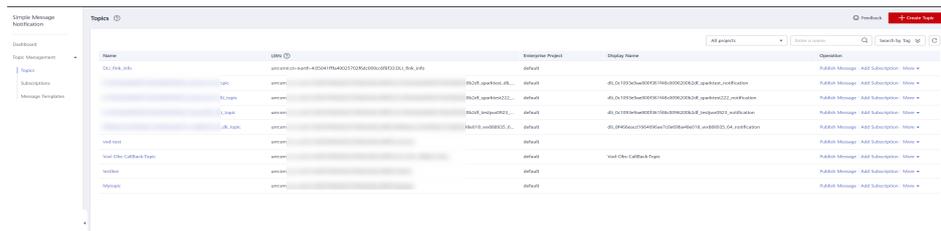
Figure 4-10 Successfully created a topic

Topic Created

011c99a26ae84a1bb963a75e7637d3fd_all_dli_topic created successfully. You can add subscription to this topic in [Topic Management](#) of SMN service.

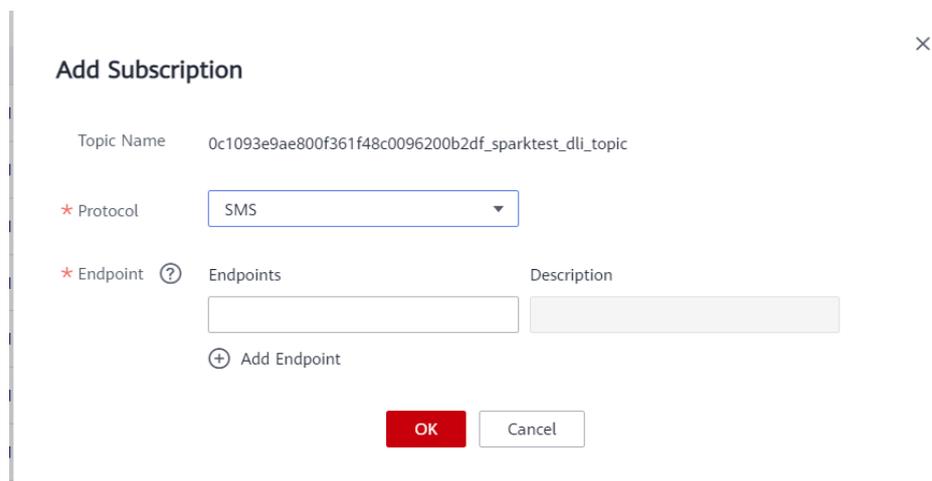
3. Click **Topic Management** in to go to the **Topic Management** page of the SMN service.

Figure 4-11 Topic management page



4. In the **Operation** column of the topic, click **Add Subscription**. Select **Protocol** to determine the subscription mode.
 - If you select **SMS** for the subscription protocol, you need to enter the mobile number for receiving the confirmation SMS message in the **Endpoint** text box.
 - If you select **Email**, enter the email address for receiving the confirmation email in the **Endpoint** text box.

Figure 4-12 Adding a subscription



5. After you click the link in the SMS message or email, you will receive a message indicating that the subscription is successful.
6. Go to the **Subscriptions** page of SMN, and check that subscription status is **Confirmed**.

4.11 Managing Queue Tags

Tag Management

A tag is a key-value pair that you can customize to identify cloud resources. It helps you to classify and search for cloud resources. A tag consists of a tag key and a tag value.

If you use tags in other cloud services, you are advised to create the same tag (key-value pairs) for cloud resources used by the same business to keep consistency.

DLI supports the following two types of tags:

- Resource tags: non-global tags created on DLI
- Predefined tags: global tags created on Tag Management Service (TMS).
For more information about predefined tags, see [Tag Management Service User Guide](#).

DLI allows you to add, modify, or delete tags for queues.

- Step 1** In the navigation pane of the DLI management console, choose **Resources > Queue Management**.
- Step 2** In the **Operation** column of the queue, choose **More > Tags**.
- Step 3** The tag management page is displayed, showing the tag information about the current queue.
- Step 4** Click **Add/Edit Tag** to switch to the **Add/Edit Tag** dialog box. Enter a tag and a value, and click **Add**.

Figure 4-13 Adding/Editing tags

Add/Edit Tag ×

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#)

To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

Table 4-6 Tag parameters

Parameter	Description
Tag key	<p>You can specify the tag key in either of the following ways:</p> <ul style="list-style-type: none">Click the text box and select a predefined tag key from the drop-down list. To add a predefined tag, you need to create one on TMS and then select it from the Tag key drop-down list. You can click View predefined tags to go to the Predefined Tags page of the TMS console. Then, click Create Tag in the upper corner of the page to create a predefined tag. For details, see Creating Predefined Tags in <i>Tag Management Service User Guide</i>.Enter a tag key in the text box. <p>NOTE A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:+=-@) are allowed, but the value cannot start or end with a space or start with _sys_.</p>
Tag value	<p>You can specify the tag value in either of the following ways:</p> <ul style="list-style-type: none">Click the text box and select a predefined tag value from the drop-down list.Enter a tag value in the text box. <p>NOTE A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:+=-@) are allowed. The value cannot start or end with a space.</p>

NOTE

- A maximum of 20 tags can be added.
- Only one tag value can be added to a tag key.
- The key name in each resource must be unique.

Step 5 Click **OK**.

Step 6 (Optional) To delete a tag, locate the row where the tag resides in the tag list and click **Delete** in the **Operation** column to delete the tag.

----End

4.12 Setting Queue Properties

Scenario

DLI allows you to set properties for queues.

You can set Spark driver parameters to improve the scheduling efficiency of queues.

This section describes how to set queue properties on the management console.

Constraints and Limitations

- Only SQL queues of the Spark engine support configuring queue properties.
- Setting queue properties is only supported after the queue has been created.
- Currently, only queue properties related to the Spark driver can be set.
- Queue properties cannot be set in batches.

Procedure

1. In the navigation pane of the DLI management console, choose **Resources > Queue Management**.
2. In the **Operation** column of the queue, choose **More > Set Property**.
3. Go to the queue property setting page and set property parameters. For details about the property parameters, see [Table 4-7](#).

Table 4-7 Queue properties

Property	Description	Value Range
Max. Spark Driver Instances	Maximum number of Spark drivers can be started on this queue, including the Spark driver that is prestarted and the Spark driver that runs jobs.	<ul style="list-style-type: none">• For a 16-CU queue, the value is 2.• For a queue that has more than 16 CUs, the value range is [2, queue CUs/16].
Max. Prestart Spark Driver Instances	Maximum number of Spark drivers can be prestarted on this queue. When the number of Spark drivers that run jobs exceeds the value of Max. Concurrency per Instance , the jobs are allocated to the Spark drivers that are prestarted.	<ul style="list-style-type: none">• For a 16-CU queue, the value range is 0 to 1.• For a queue that has more than 16 CUs, the value range is [2, queue CUs/16].
Max. Concurrency per Instance	Maximum number of jobs can be concurrently executed by a Spark driver. When the number of jobs exceeds the value of this parameter, the jobs are allocated to other Spark drivers.	1-32

4. Click **OK**.

5 Data Management

5.1 Databases and Tables

5.1.1 Overview

DLI database and table management provide the following functions:

- [Database Permission Management](#)
- [Table Permission Management](#)
- [Creating a Database or a Table](#)
- [Deleting a Database or a Table](#)
- [Changing the Owners of Databases and Tables](#)
- [Importing Data](#)
- [Exporting Data](#)
- [Viewing Metadata](#)
- [Previewing Data](#)

Difference Between DLI Tables and OBS Tables

- Data stored in DLI tables is applicable to delay-sensitive services, such as interactive queries.
- Data stored in OBS tables is applicable to delay-insensitive services, such as historical data statistics and analysis.

Constraints

- **Database**
 - **default** is the database built in DLI. You cannot create a database named **default**.
 - DLI supports a maximum of 50 databases.
- **Table**
 - DLI supports a maximum of 5,000 tables.

- DLI supports the following table types:
 - **MANAGED**: Data is stored in a DLI table.
 - **EXTERNAL**: Data is stored in an OBS table.
 - **View**: A view can only be created using SQL statements.
 - **Datasource table**: The table type is also **EXTERNAL**.
- You cannot specify a storage path when creating a DLI table.
- **Data import**
 - Only OBS data can be imported to DLI or OBS.
 - You can import data in CSV, Parquet, ORC, JSON, or Avro format from OBS to tables created on DLI.
 - To import data in CSV format to a partitioned table, place the partition column in the last column of the data source.
 - The encoding format of imported data can only be UTF-8.
- **Data export**
 - Data in DLI tables (whose table type is **MANAGED**) can only be exported to OBS buckets, and the export path must contain a folder.
 - The exported file is in JSON format, and the text format can only be UTF-8.
 - Data can be exported across accounts. That is, after account B authorizes account A, account A has the permission to read the metadata and permission information of account B's OBS bucket as well as the read and write permissions on the path. Account A can export data to the OBS path of account B.

Databases and Tables Page

The **Databases and Tables** page displays all created databases. You can view the database information, such as the owner and the number of tables.

Table 5-1 Database and table management parameters

Parameter	Description
Database Name	<ul style="list-style-type: none">● The database name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_).● The database name is case insensitive and cannot be left unspecified.● It cannot exceed 128 characters.
Username	Database owner.
Tables	Number of tables in the database.
Description	Description of the database specified during database creation. If no description is provided, -- is displayed.

Parameter	Description
Enterprise Project	Enterprise project to which the database belongs. An enterprise project facilitates project-level management and grouping of cloud resources and users. For details about how to set enterprise projects, see Enterprise Management User Guide .
Operation	<ul style="list-style-type: none">● Permissions: View the permission information and perform user authorization, permission settings, and user permission revocation.● Tables: View the tables in the corresponding database. For details, see Table Management Page.● Create Table: This permission allows you to create a table in the corresponding database.● Modify Database. This permission allows you to change the owner of the database. The username must exist under the same account.● Drop Database: This permission allows you to delete the selected database.● Modify Enterprise Project: Add the enterprise project.● Tags: Add or edit tags.

Table Management Page

From the **Data Management** page, click the database name or **Tables** in the **Operation** column to switch to the table management page.

The displayed page lists all tables created in the current database. You can view the table type, data storage location, and other information. Tables are listed in chronological order by default, with the most recently created tables displayed at the top.

Table 5-2 Table management parameters

Parameter	Description
Table Name	<ul style="list-style-type: none">● The table name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_).● The table name is case insensitive and cannot be left unspecified.● The table name can contain the dollar sign (\$). An example value is \$test.● It cannot exceed 128 characters.

Parameter	Description
Table Type	Table type. Available options are as follows: <ul style="list-style-type: none">● Managed: Indicates that data is stored in a DLI table.● External: Indicates that data is stored in an OBS table.● View: Indicates the view type. You can only create views using SQL statements. NOTE The table or view information contained in the view cannot be modified. If the table or view information is modified, the query may fail.
Owner	User who creates the table.
Storage Location	DLI, OBS, View, CloudTable, and CSS data location
Size	Size of the data in the table. The value is displayed only for tables of the Managed type. For tables of other types, -- is displayed.
Data Source Path	<ul style="list-style-type: none">● If Data Location is OBS, the corresponding OBS path is displayed.● If Data Location is DLI and View, -- is displayed.● When the data storage location is a datasource connection service such as CloudTable and CSS, the corresponding URL is displayed.
Created	Time when the table is created.
Last Accessed	Last time when an operation was performed on the table.
Operation	<ul style="list-style-type: none">● Manage Permissions: This operation allows you to view the permission information and perform user authorization, permission settings, and user permission revocation.● More:<ul style="list-style-type: none">– Delete: Delete a table from the corresponding database.– Modify Owner: Change the owner of a table The username must exist under the same account.– Import: Import data stored in an OBS bucket to a DLI or OBS table.– Properties: View data in Metadata and Preview tabs.

5.1.2 Managing Database Permissions

Scenario

- You can isolate databases allocated to different users by setting permissions to ensure data query performance.

- The administrator and database owner have all permissions, which cannot be set or modified by other users.
- When setting database permissions for a new user, ensure that the region of the user group to which the user belongs has the **Tenant Guest** permission. For details about the Tenant Guest permission and how to apply for the permission, see [Creating a User Group and Assigning Permissions](#) and [System Permissions](#).

Precautions

- Lower-level objects automatically inherit permissions granted to upper-level objects. The hierarchical relationship is database > table > column.
- The database owner, table owner, and **authorized** users can assign permissions on the database and tables.
- Columns can only inherit the query permission. For details about **Inheritable Permissions**, see [Managing Database Permissions](#).
- The permissions can be revoked only at the initial level to which the permissions are granted. You need to grant and revoke permissions at the same level. For example, after you are granted the insertion permission on a database, you can obtain the insertion permission on the tables in the database. Your insertion permission can be revoked only at the database level.
- If you create a database with the same name as a deleted database, the database permissions will not be inherited. In this case, you need to grant the database permissions to users or projects.

For example, user A is granted with the permission to delete the **testdb** database. Delete the database and create another one with the same name. You need to grant user A the deletion permission of the **testdb** database again.

Viewing Database Permissions

1. On the left of the management console, choose **Data Management > Databases and Tables**.
2. Locate the row where the target database resides and click **Manage Permissions** in the **Operation** column.

Permissions can be granted to new users or projects, modified for users or projects with existing permissions, or revoked from a user or project.

Granting Permissions to a New User or Project

Here, the new user or project refers to a user or a project that does not have permissions on the database.

1. Click a database you need. In the displayed **Database Permission Management** page, click **Grant Permission** in the upper right corner.
2. In the displayed dialog box, select **User** or **Project**, enter the username or select the project to be authorized, and select the required permissions. For details about the permissions, see [Table 5-3](#).

Table 5-3 Parameters

Parameter	Description
Authorization Object	Select User or Project .
Username/ Project Name	<ul style="list-style-type: none"> • If you select User, enter the IAM username when adding a user to the database. <p>NOTE The username is an existing IAM user name and has logged in to the DLI management console.</p> <ul style="list-style-type: none"> • If you select Project, select the project to be authorized in the current region. <p>NOTE When Project is selected:</p> <ul style="list-style-type: none"> • If you set Non-inheritable Permissions, you cannot view tables in the corresponding database within the project. • If you set Inheritable Permissions, you can view all tables in the database within the project.

Parameter	Description
Non-Inherited Permissions	<p>Select a permission to grant it to the user, or deselect a permission to revoke it.</p> <p>Non-inherited permissions apply only to the current database.</p> <ul style="list-style-type: none">• The following permissions are applicable to both user and project authorization:<ul style="list-style-type: none">– Drop Database: This permission allows you to delete the current database.– Create Table: This permission allows you to create tables in the current database.– Create View: This permission allows you to create views in the current database.– Execute SQL EXPLAIN: This permission allows you to execute an EXPLAIN statement and view information about how this database executes a query.– Create Role: This permission allows you to create roles in the current database.– Delete Role: This permission allows you to delete roles of the current database.– View Role: This permission allows you to view the role of the current user.– Bind Role: This permission allows you to bind roles to the current database.– Unbind Role: This permission allows you to bind roles from the current database.– View All Binding Relationships: This permission allows you to view the binding relationships between all roles and users.– Create Function: This permission allows you to create a function in the current database.– Delete Function: This permission allows you to delete functions from the current database.– View All Functions: This permission allows you to view all functions in the current database.– View Function Details: This permission allows you to view details about the current function.• The following permissions can only be granted to users:<ul style="list-style-type: none">– View All Tables: This permission allows you to view all tables in the current database. <p>NOTE</p> <p>If this permission of a specific database is not granted, all tables in the database will not be displayed.</p> <ul style="list-style-type: none">– View Database: This permission allows you to view the information about the current database.

Parameter	Description
	NOTE If this permission is not granted, the database will not be displayed.

Parameter	Description
Inherited Permissions	<p>Select a permission to grant it to the user, or deselect a permission to revoke it.</p> <p>Inherited permissions are applicable to the current database and all its tables. However, only the query permission is applicable to table columns.</p> <p>The following permissions can be granted to both user and project.</p> <ul style="list-style-type: none">● Drop Table: This permission allows you to delete tables in a database.● Select Table: This permission allows you to query data of the current table.● View Table Information: This permission allows you to view information about the current table.● Insert: This permission allows you to insert data into the current table.● Add Column: This permission allows you to add columns to the current table.● Overwrite: This permission allows you to insert data to overwrite the data in the current table.● Grant Permission: This permission allows you to grant database permissions to other users or projects.● Revoke Permission: This permission allows you to revoke the permissions of the database that other users have but cannot revoke the database owner's permissions.● Add Partition to Partition Table: This permission allows you to add a partition to a partition table.● Delete Partition from Partition Table: This permission allows you to delete existing partitions from a partition table.● Configure Path for Partition: This permission allows you to set the path of a partition in a partition table to a specified OBS path.● Rename Table Partition: This permission allows you to rename partitions in a partition table.● Rename Table: This permission allows you to rename tables.● Restore Table Partition: This permission allows you to export partition information from the file system and save the information to metadata.● View All Partitions: This permission allows you to view all partitions in a partition table.● View Other Users' Permissions: This permission allows you to query other users' permission on the current database.

3. Click **OK**.

Modifying Permissions for an Existing User or Project

For a user or project that has some permissions on the database, you can revoke the existing permissions or grant new ones.

NOTE

If the options in **Set Permission** are gray, the corresponding account does not have the permission to modify the database. You can apply to the administrator, database owner, or other authorized users for granting and revoking permissions of databases.

1. In the **User Permission Info** list, find the user whose permission needs to be set.
 - If the user is a sub-user, you can set permissions for it.
 - If the user is already an administrator, you can only view the permissions information.

In the **Project Permission Info** list, locate the project for which you want to set permissions and click **Set Permission**.

2. In the **Operation** column of the sub-user or project, click **Set Permission**. The **Set Permission** dialog box is displayed.

For details about the permissions of database users or projects, see [Table 5-3](#).

3. Click **OK**.

Revoking All Permissions of a User or Project

Revoke all permissions of a user or a project.

- In the user list under **User Permission Info**, locate the row where the target sub-user resides and click **Revoke Permission** in the **Operation** column. In the displayed dialog box, click **OK**. In this case, the user has no permissions on the database.

NOTE

If a user is an administrator, **Revoke Permission** is gray, indicating that the user's permission cannot be revoked.

- In the **Project Permission Info** area, select the project whose permissions need to be revoked and click **Revoke Permission** in the **Operation** column. After you click **OK**, the project does not have any permissions on the database.

5.1.3 Managing Table Permissions

Operation Scenario

- You can isolate databases allocated to different users by setting permissions to ensure data query performance.
- The administrator and database owner have all permissions, which cannot be set or modified by other users.

- When setting database permissions for a new user, ensure that the user group to which the user belongs has the **Tenant Guest** permission.

Precautions

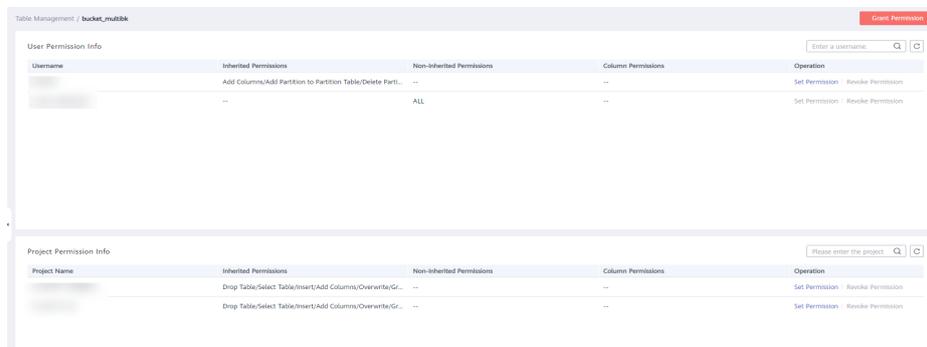
- By the rules in [Common Operations Supported by DLI System Policy](#), you need to authorize a user to view tables in a database of the owner account.
- If you create a table with the same name as a deleted table, the table permissions will not be inherited. In this case, you need to grant the table permissions to users or projects.

For example, user A is granted with the permission to delete the **testTable** table. Delete the table and create another one with the same name. You need to grant user A the deletion permission of the **testTable** table again.

Viewing Table Permissions

1. On the left of the management console, choose **Data Management > Databases and Tables**.
2. Click the database name in the table whose authority is to be set. The **Table Management** page of the database is displayed.
3. Locate the row where the target table resides and click **Manage Permissions** in the **Operation** column.

Figure 5-1 Table permission management



Permissions can be granted to new users or projects, modified for users or projects with existing permissions, or revoked from a user or project.

Granting Permissions to a New User or a Project

Here, the new user or project refers to a user or a project that does not have permissions on the database.

1. Click the table you need. In the displayed table permissions page, click **Grant Permission** in the upper right corner.
2. In the displayed **Grant Permission** dialog box, select the required permissions.
 - For details about the DLI table permissions, see [Table 5-4](#).

Figure 5-2 Granting DLI table permissions to a user

Grant Permission ×

* Authorization Object: User Project

* Username:

Select the permissions to be granted to the user

Non-inheritable Permissions

<input type="checkbox"/> Select all	<input type="checkbox"/> Select Table	<input type="checkbox"/> View Table Info
<input type="checkbox"/> View Table	<input type="checkbox"/> Drop Table	<input type="checkbox"/> Rename Table
<input type="checkbox"/> View Table Creation Statement	<input type="checkbox"/> Overwrite Table	<input type="checkbox"/> Add Column
<input type="checkbox"/> Insert	<input type="checkbox"/> View Other User's Permissions	<input type="checkbox"/> Grant Permission
<input type="checkbox"/> Access Metadata		
<input type="checkbox"/> Revoke Permission		

OK Cancel

Figure 5-3 Granting DLI table permissions to a project

Grant Permission ×

* Authorization Object: User Project

* Project:

Select the permissions to be granted to the user

Non-inheritable Permissions

<input type="checkbox"/> Select all	<input type="checkbox"/> View Table Info	<input type="checkbox"/> View Table Creation Statement
<input type="checkbox"/> Select Table	<input type="checkbox"/> Rename Table	<input type="checkbox"/> Insert
<input type="checkbox"/> Drop Table	<input type="checkbox"/> Add Column	<input type="checkbox"/> Access Metadata
<input type="checkbox"/> Overwrite Table	<input type="checkbox"/> Grant Permission	<input type="checkbox"/> Revoke Permission
<input type="checkbox"/> View Other User's Permissions		

OK Cancel

Table 5-4 Parameter description

Parameter	Description
Authorization Object	Select User or Project .
Username/Project	<ul style="list-style-type: none"> If you select User, enter the IAM username when granting table permissions to the user. <p>NOTE The username is an existing IAM user name and has logged in to the DLI management console.</p> <ul style="list-style-type: none"> If you select Project, select the project to be authorized in the current region. <p>NOTE If you select Project, you can only view information about the authorized tables and their databases.</p>

Parameter	Description
Non-inheritable Permissions	<p>Select a permission to grant it to the user, or deselect a permission to revoke it.</p> <ul style="list-style-type: none"> ● The following permissions are applicable to both user and project authorization: <ul style="list-style-type: none"> – Select Table: This permission allows you to query data of the current table. – View Table Information: This permission allows you to view information about the current table. – View Table Creation Statement: This permission allows you to view the statement for creating the current table. – Drop Table: This permission allows you to delete the current table. – Rename Table: Rename the current table. – Insert: This permission allows you to insert data into the current table. – Overwrite: This permission allows you to insert data to overwrite the data in the current table. – Add Column: Add columns to the current table. – Grant Permission: The current user can grant table permissions to other users. – Revoke Permission: The current user can revoke the table's permissions that other users have but cannot revoke the table owner's permissions. – View Other Users' Permissions: This permission allows you to query other users' permission on the current table. <p>The partition table also has the following permissions:</p> <ul style="list-style-type: none"> – Delete Partition: This permission allows you to delete existing partitions from a partition table. – View All Partitions: This permission allows you to view all partitions in a partition table. <ul style="list-style-type: none"> ● The following permissions can only be granted to users: <ul style="list-style-type: none"> – View Table: This permission allows you to display the current table.

- For details about the OBS table permissions, see [Table 5-5](#).

Figure 5-4 Granting OBS table permissions to a user

Figure 5-5 Granting OBS table permissions to a project

Table 5-5 Parameter description

Parameter	Description
Authorization Object	Select User or Project .
Username/ Project	<ul style="list-style-type: none"> If you select User, enter the IAM username when granting table permissions to the user. NOTE The username is an existing IAM user name and has logged in to the DLI management console. If you select Project, select the project to be authorized in the current region. NOTE If you select Project, you can only view information about the authorized tables and their databases.

Parameter	Description
Non-inheritable Permissions	<p>Select a permission to grant it to the user, or deselect a permission to revoke it.</p> <ul style="list-style-type: none"> • The following permissions are applicable to both user and project authorization: <ul style="list-style-type: none"> - View Table Creation Statement: This permission allows you to view the statement for creating the current table. - View Table Information: This permission allows you to view information about the current table. - Select Table: This permission allows you to query data of the current table. - Drop Table: This permission allows you to delete the current table. - Rename Table: Rename the current table. - Insert: This permission allows you to insert data into the current table. - Overwrite: This permission allows you to insert data to overwrite the data in the current table. - Add Column: This permission allows you to add columns to the current table. - Grant Permission: This permission allows you to grant table permissions to other users or projects. - Revoke Permission: This permission allows you to revoke the table's permissions that other users or projects have but cannot revoke the table owner's permissions. - View Other Users' Permissions: This permission allows you to query other users' permission on the current table. <p>The partition table also has the following permissions:</p> <ul style="list-style-type: none"> - Add Partition: This permission allows you to add a partition to a partition table. - Delete Partition: This permission allows you to delete existing partitions from a partition table. - Configure Path for Partition: This permission allows you to set the path of a partition in a partition table to a specified OBS path. - Rename Table Partition: This permission allows you to rename partitions in a partition table. - Restore Table Partition: This permission allows you to export partition information from the file system and save the information to metadata. - View All Partitions: This permission allows you to view all partitions in a partition table.

Parameter	Description
	<ul style="list-style-type: none"> The following permissions can only be granted to users: <ul style="list-style-type: none"> View Table: This permission allows you to view the current table.

- For details about the view permissions, see [Table 5-6](#).

 **NOTE**

A view can be created only by using SQL statements. You cannot create a view on the **Create Table** page.

Figure 5-6 Granting view permissions to a user

Grant Permission

* Authorization Object User Project

* Username

Select the permissions to be granted to the user

Non-Inherited Permissions

<input type="checkbox"/> Select all	<input type="checkbox"/> View Table Info	<input type="checkbox"/> View Table Creation Statement
<input type="checkbox"/> View Table	<input type="checkbox"/> Select Table	<input type="checkbox"/> Rename Table
<input type="checkbox"/> Drop Table	<input type="checkbox"/> Grant Permission	<input type="checkbox"/> Revoke Permission
<input type="checkbox"/> View Other User's Permissions		

Figure 5-7 Granting view permissions to a project

Grant Permission

* Authorization Object User Project

* Enterprise Project

Select the permissions to be granted to the user

Non-Inherited Permissions

<input type="checkbox"/> Select all	<input type="checkbox"/> View Table Creation Statement	<input type="checkbox"/> Drop Table
<input type="checkbox"/> View Table Info	<input type="checkbox"/> Rename Table	<input type="checkbox"/> View Other User's Permissions
<input type="checkbox"/> Select Table	<input type="checkbox"/> Revoke Permission	
<input type="checkbox"/> Grant Permission		

Table 5-6 Parameter description

Parameter	Description
Authorization Object	Select User or Project .

Parameter	Description
Username/ Project	<ul style="list-style-type: none"> ● If you select User, enter the IAM username when adding a user to the database. <p>NOTE The username is an existing IAM user name and has logged in to the DLI management console.</p> <ul style="list-style-type: none"> ● If you select Project, select the project to be authorized in the current region. <p>NOTE If you select Project, you can only view information about the authorized tables and their databases.</p>
Non-inheritable Permissions	<p>Select a permission to grant it to the user, or deselect a permission to revoke it.</p> <ul style="list-style-type: none"> ● The following permissions are applicable to both user and project authorization: <ul style="list-style-type: none"> - View Table Information: This permission allows you to view information about the current table. - View Table Creation Statement: This permission allows you to view the statement for creating the current table. - Drop Table: This permission allows you to delete the current table. - Select Table: This permission allows you to query data of the current table. - Rename Table: Rename the current table. - Grant Permission: The current user or project can grant table permissions to other users or projects. - Revoke Permission: The current user or project can revoke the table's permissions that other users or projects have but cannot revoke the table owner's permissions. - View Other Users' Permissions: This permission allows you to query other users' permission on the current table. ● Only applicable to <ul style="list-style-type: none"> - View Table: This permission allows you to view the current table.

3. Click **OK**.

Modifying Permissions for an Existing User or Project

For a user or project that has some permissions on the database, you can revoke the existing permissions or grant new ones.

 **NOTE**

If all options under **Set Permission** are gray, you are not allowed to change permissions on this table. You can apply to the administrator, table owner, or other authorized users for granting and revoking table permissions.

1. In the **User Permission Info** list, find the user whose permission needs to be set.
 - If the user is a sub-user and is not the owner of the table, you can set permissions.
 - If the user is an administrator or table owner, you can only view permissions.

In the **Project Permission Info** list, locate the project for which you want to set permissions and click **Set Permission**.

2. In the **Operation** column of the sub-user or project, click **Set Permission**. The **Set Permission** dialog box is displayed.
 - For details about DLI table user or project permissions, see [Table 5-4](#).
 - For details about OBS table user or project permissions, see [Table 5-5](#).
 - For details about View table user or project permissions, see [Table 5-6](#).
3. Click **OK**.

Revoking All Permissions of a User or Project

Revoke all permissions of a user or a project.

- In the user list under **User Permission Info**, locate the row where the target sub-user resides and click **Revoke Permission** in the **Operation** column. In the displayed dialog box, click **OK**. In this case, the user has no permissions on the table.

 **NOTE**

In the following cases, **Revoke Permission** is gray, indicating that the permission of the user cannot be revoked.

- The user is an administrator.
 - The sub-user is the owner of the table.
 - The sub-user has only inheritable permissions.
- In the **Project Permission Info** area, select the project whose permissions need to be revoked and click **Revoke Permission** in the **Operation** column. After you click **OK**, the project does not have any permissions on the table.

 **NOTE**

If a project has only inheritable permissions, **Revoke Permission** is gray, indicating that the permissions of the project cannot be revoked.

5.1.4 Creating a Database or a Table

Definition of Database and Table in DLI

A database, built on the computer storage device, is a data warehouse where data is organized, stored, and managed based on its structure.

The table is an important part of the database. It consists of rows and columns. Each column functions as a field. Each value in a field (column) represents a type of data.

The database is a framework and the table contains data content. A database has one or more tables.

You can create databases and tables on the management console or using SQL statements. This section describes how to create a database and a table on the management console.

For details about using SQL statements, see [Creating a Database](#), [Creating an OBS Table](#), and [Creating a DLI Table](#).

NOTE

A view can be created only by using SQL statements. You cannot create a view on the **Create Table** page.

Precautions

- If a folder and a file have the same name in the OBS directory, the file path is preferred as the path of the OBS table to be created.
- When creating a database or table, you need to grant permissions to other users so that they can view the created database or table. For details, see [Common Operations Supported by DLI System Policy](#).

Creating a Database

Step 1 You can create a database on either the **Data Management** page or the **SQL Editor** page.

- To create a database on the **Data Management** page:
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. In the upper right corner of the **Databases and Tables** page, click **Create Database** to create a database.
- To create a database on the **SQL Editor** page:
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the navigation pane on the left, click  beside **Databases**.

Step 2 In the displayed **Create Database** dialog box, specify **Name** and **Description** by referring to [Table 5-7](#).

Figure 5-8 Creating a database

Create Database

×

You can create 1 more databases. [Increase quota.](#)

★ Databases

Description

0/128

★ Enterprise Project Enterprise Project ↻ ⓘ Create Enterprise Project

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#) ↻

To add a tag, enter a tag key and a tag value below.

Tags

10 tags available for addition.

Table 5-7 Description

Parameter	Description
Database Name	<ul style="list-style-type: none"> The database name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_). The database name is case insensitive and cannot be left blank. The length of the database name cannot exceed 128 characters. <p>NOTE The default database is a built-in database. You cannot create the default database.</p>
Enterprise Project	<p>If the created queue belongs to an enterprise project, you can select the corresponding enterprise project.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users.</p> <p>For details about how to set enterprise projects, see Enterprise Management User Guide.</p> <p>NOTE This parameter is displayed only for users who have enabled the Enterprise Management Service.</p>
Description	Description of a database.

Parameter	Description
Tag	<p>Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS).</p> <p>For details, see Tag Management Service User Guide.</p> <p>NOTE</p> <ul style="list-style-type: none">• A maximum of 20 tags can be added.• Only one tag value can be added to a tag key.• The key name in each resource must be unique.• Tag key: Enter a tag key name in the text box. <p>NOTE</p> <p>A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed, but the value cannot start or end with a space or start with <code>_sys_</code>.</p> <ul style="list-style-type: none">• Tag value: Enter a tag value in the text box. <p>NOTE</p> <p>A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed. The value cannot start or end with a space.</p>

Step 3 Click **OK**.

After a database is created, you can view and select the database for use on the **Databases and Tables** page or **SQL Editor** page.

----End

Creating a Table

Before creating a table, ensure that a database has been created.

Step 1 You can create a table on either the **Databases and Tables** page or the **SQL Editor** page.

NOTE

Datasource connection tables, such as View tables, HBase (CloudTable/MRS) tables, OpenTSDB (CloudTable/MRS) tables, DWS tables, RDS tables, and CSS tables, cannot be created. You can use SQL to create views and datasource connection tables. For details, see sections **Creating a View** and **Creating a Datasource Connection Table** in the *Data Lake Insight SQL Syntax Reference*.

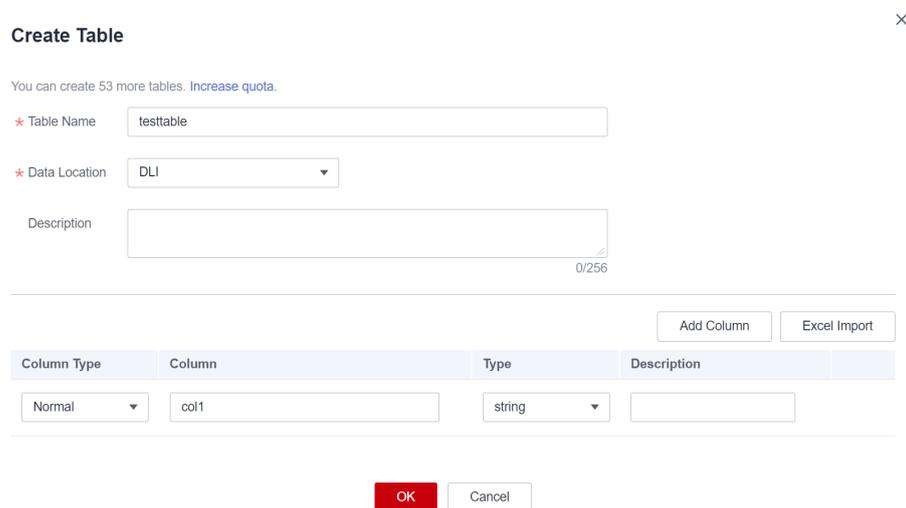
- To create a table on the **Data Management** page:
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. On the **Databases and Tables** page, select the database for which you want to create a table. In the **Operation** column, click **More > Create Table** to create a table in the current database.

- To create a table on the **SQL Editor** page:
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the navigation pane of the displayed **SQL Editor** page, click **Databases**. You can create a table in either of the following ways:
 - Click a database name. In the **Tables** area, click  on the right to create a table in the current database.
 - Click  on the right of the database and choose **Create Table** from the shortcut menu to create a table in the current database.

Step 2 In the displayed **Create Table** dialog box, set parameters as required.

- If you set **Data Location** to **DLI**, set related parameters by referring to [Table 5-8](#).

Figure 5-9 Creating a DLI table



Create Table ×

You can create 53 more tables. [Increase quota](#).

* Table Name

* Data Location

Description 0/256

Column Type	Column	Type	Description
<input type="text" value="Normal"/>	<input type="text" value="col1"/>	<input type="text" value="string"/>	<input type="text"/>

- If you set **Data Location** to **OBS**, set related parameters by referring to [Table 5-8](#) and [Table 5-9](#).

Figure 5-10 Creating an OBS table

Create Table
×

You can create 53 more tables. [Increase quota.](#)

* Table Name

* Data Location

* Data Format

* Storage Path

Description

Column Type	Column	Type	Description
Normal	col1	string	

Advanced Settings

Table 5-8 Common parameters

Parameter	Description	Example
Table Name	<ul style="list-style-type: none"> - The table name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_). - The table name is case insensitive and cannot be left unspecified. - The table name can contain the dollar sign (\$). An example value is \$test. - The length of the table name cannot exceed 128 characters. 	table01
Data Location	Data storage location. Currently, DLI and OBS are supported.	DLI
Description	Description of the table.	-
Column Type	Available values: Normal or Partition	Normal

Parameter	Description	Example
Column	<p>Name of a column in a table. The column name must contain at least one letter and can contain underscores (_). It cannot contain only digits.</p> <p>You can select Normal or Partition. Partition columns are dedicated to partition tables. User data is partitioned to improve query efficiency.</p> <p>NOTE The column name is case-insensitive and must be unique.</p>	name
Type	<p>Data type of a column. This parameter corresponds to Column Name.</p> <ul style="list-style-type: none"> - string: The data is of the string type. - int: Each integer is stored on four bytes. - date: The value ranges from 0000-01-01 to 9999-12-31. - double: Each number is stored on eight bytes. - boolean: Each value is stored on one byte. - decimal: The valid bits are positive integers between 1 to 38, including 1 and 38. The decimal digits are integers less than 10. - smallint/short: The number is stored on two bytes. - bigint/long: The number is stored on eight bytes. - timestamp: The data indicates a date and time. The value can be accurate to six decimal points. - float: Each number is stored on four bytes. - tinyint: Each number is stored on one byte. Only OBS tables support this data type. 	string
Column Description	Description of a column.	-
Operation	<ul style="list-style-type: none"> - Add Column - Delete <p>NOTE If the table to be created includes a great number of columns, you are advised to use SQL statements to create the table or import column information from the local EXCEL file.</p>	-

Table 5-9 Parameter description when **Data Location** is set to **OBS**

Parameter	Description	Example
Data Format	<p>DLI supports the following data formats:</p> <ul style="list-style-type: none"> - Parquet: DLI can read non-compressed data or data that is compressed using Snappy and gzip. - CSV: DLI can read non-compressed data or data that is compressed using gzip. - ORC: DLI can read non-compressed data or data that is compressed using Snappy. - JSON: DLI can read non-compressed data or data that is compressed using gzip. - Avro: DLI can read uncompressed Avro data. 	CSV
Storage Path	<p>Enter or select an OBS path. The path can be a folder or a path.</p> <p>NOTE If you need to import data stored in OBS to the OBS table, set this parameter to the path of a folder. If the table creation path is a file, data fails to be imported.</p>	obs://obs1/ sampledata.csv
Table Header: No/Yes	<p>This parameter is valid only when Data Format is set to CSV. Whether the data source to be imported contains the table header.</p> <p>Click Advanced Settings and select the check box next to Table Header: No. If the check box is selected, the table header is displayed. If the check box is deselected, no table header is displayed.</p>	-
User-defined Delimiter	<p>This parameter is valid only when Data Format is set to CSV and you select User-defined Delimiter.</p> <p>The following delimiters are supported:</p> <ul style="list-style-type: none"> - Comma (,) - Vertical bar () - Tab character (\t) - Others: Enter a user-defined delimiter. 	Comma (,)

Parameter	Description	Example
User-defined Quotation Character	This parameter is valid only when Data Format is set to CSV and you select User-defined Quotation Character . The following quotation characters are supported: <ul style="list-style-type: none"> – Single quotation mark (') – Double quotation marks (") – Others: Enter a user-defined quotation character. 	Single quotation mark (')
User-defined Escape Character	This parameter is valid only when Data Format is set to CSV and you select User-defined Escape Character . The following escape characters are supported: <ul style="list-style-type: none"> – Backslash (\) – Others: Enter a user-defined escape character. 	Backslash (\)
Date Format	This parameter is valid only when Data Format is set to CSV or JSON . This parameter specifies the format of the date in the table and is valid only Advanced Settings is selected. The default value is yyyy-MM-dd . For definition of characters involved in the date pattern, see Table 3 in Importing Data to the Table .	2000-01-01
Timestamp Format	This parameter is valid only when Data Format is set to CSV or JSON . This parameter specifies the format of the timestamp in the table and is valid only Advanced Settings is selected. The default value is yyyy-MM-dd HH:mm:ss . For definition of characters involved in the time pattern, see Table 3 in Importing Data to the Table .	2000-01-01 09:00:00

Step 3 Click **OK**.

After a table is created, you can view and select the table for use on the **Data Management** page or **SQL Editor** page.

Step 4 (Optional) After a DLI table is created, you can decide whether to directly import data to the table.

----End

5.1.5 Deleting a Database or a Table

You can delete unnecessary databases and tables based on actual conditions.

Precautions

- You are not allowed to delete databases or tables that are being used for running jobs.
- The administrator, database owner, and users with the database deletion permission can delete the database. The administrator, database owner, and users with the table deletion permission can delete the table.

NOTE

If a database or table is deleted, it cannot be recovered. Exercise caution when performing this operation.

Deleting a Table

You can delete a table on either the **Data Management** page or the **SQL Editor** page.

- Delete the table on the **Data Management** page.
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. Locate the row where the database whose tables you want to delete, click the database name to switch to the **Table Management** page.
 - c. Locate the row where the target table locates and click **More > Delete** in the **Operation** column.
 - d. In the displayed dialog box, click **Yes**.
- Delete a table on the **SQL Editor** page.
 - a. On the top menu bar of the DLI management console, click **SQL Editor**.
 - b. In the navigation tree on the left, click **Databases**. Click the name of a database where the table you want belongs. The tables of the selected database are displayed.
 - c. Click  on the right of the table and choose **Delete** from the shortcut menu.
 - d. In the dialog box that is displayed, click **OK**.

Deleting a Database

1. On the left of the management console, choose **Data Management > Databases and Tables**.
2. Locate the row where the target database locates and click **More > Drop Database** in the **Operation** column.

NOTE

You cannot delete databases that contain tables. To delete a database containing tables, delete the tables first.

3. In the displayed dialog box, click **Yes**.

5.1.6 Modifying the Owners of Databases and Tables

During actual use, developers create databases and tables and submit them to test personnel for testing. After the test is complete, the databases and tables are transferred to O&M personnel for user experience. In this case, you can change the owner of the databases and tables to transfer data to other owners.

Modifying the Database Owner

You can change the owner of a database on either the **Data Management** page or the **SQL Editor** page.

- On the **Data Management** page, change the database owner.
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. On the **Databases and Tables** page, locate the database you want and click **More > Modify Database** in the **Operation** column.
 - c. In the displayed dialog box, enter a new owner name (an existing username) and click **OK**.
- Change the database owner on the **SQL Editor** page.
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the navigation tree on the left, click **Databases**, click  on the right of the database you want to modify, and choose **Modify Database** from the shortcut menu.
 - c. In the displayed dialog box, enter a new owner name (an existing username) and click **OK**.

Modifying the Table Owner

1. On the left of the management console, choose **Data Management > Databases and Tables**.
2. Click the name of the database corresponding to the table to be modified. The **Manage Tables** page of the database is displayed.
3. In the **Operation** column of the target table, choose **More > Modify Owner**.
4. In the displayed dialog box, enter a new owner name (an existing username) and click **OK**.

5.1.7 Importing Data to the Table

You can import data from OBS to a table created in DLI.

Precautions

- Only one path can be specified during data import. The path cannot contain commas (,).
- To import data in CSV format to a partitioned table, place the column to be partitioned in the last column of the data source.
- You are advised not to concurrently import data in to a table. If you concurrently import data into a table, there is a possibility that conflicts occur, leading to failed data import.

- The imported file can be in CSV, Parquet, ORC, JSON, and Avro format. The encoding format must be UTF-8.

Prerequisites

The data to be imported has been stored on OBS.

Procedure

Step 1 You can import data on either the **Data Management** page or the **SQL Editor** page.

- To import data on the **Data Management** page:
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. Click the name of the database corresponding to the table where data is to be imported to switch to the table management page.
 - c. Locate the row where the target table resides and choose **More > Import** in the **Operation** column. The **Import** dialog box is displayed.

Figure 5-11 Importing data

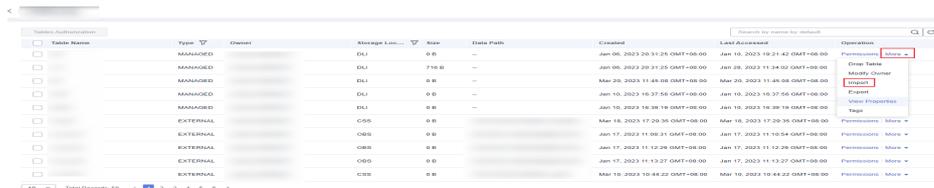
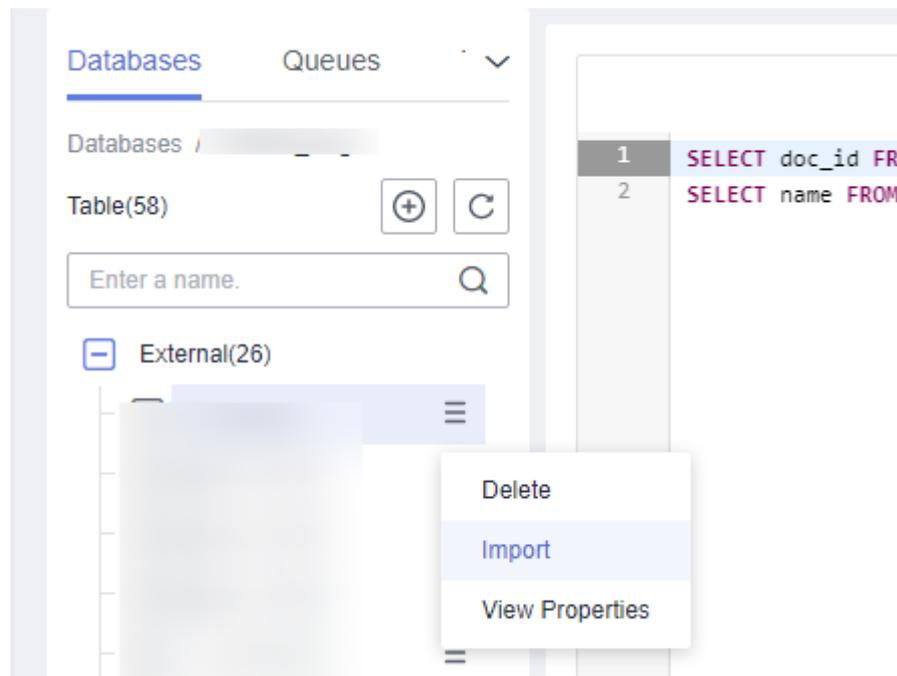


Table Name	Type	Owner	Message Size	Size	State	Created	Last Updated	Operation
	MANAGED		DLI	0 B	---	Jan 06, 2023 20:31:25 GMT+08:00	Jan 10, 2023 10:21:42 GMT+08:00	Permissions More
	MANAGED		DLI	710 B	---	Jan 06, 2023 20:31:25 GMT+08:00	Jan 28, 2023 11:24:02 GMT+08:00	Create Table
	MANAGED		DLI	0 B	---	Mar 20, 2023 11:48:08 GMT+08:00	Mar 20, 2023 11:28:08 GMT+08:00	Modify Owner
	MANAGED		DLI	0 B	---	Jan 10, 2023 10:37:55 GMT+08:00	Jan 10, 2023 10:37:55 GMT+08:00	Export
	MANAGED		DLI	0 B	---	Jan 10, 2023 10:39:18 GMT+08:00	Jan 10, 2023 10:39:18 GMT+08:00	Open Properties
	MANAGED		DLI	0 B	---	Jan 10, 2023 10:39:18 GMT+08:00	Jan 10, 2023 10:39:18 GMT+08:00	Tags
	EXTERNAL		ORC	0 B	---	Mar 16, 2023 17:20:26 GMT+08:00	Mar 16, 2023 17:20:26 GMT+08:00	Permissions More
	EXTERNAL		ORC	0 B	---	Jan 17, 2023 11:08:21 GMT+08:00	Jan 17, 2023 11:10:24 GMT+08:00	Permissions More
	EXTERNAL		ORC	0 B	---	Jan 17, 2023 11:12:28 GMT+08:00	Jan 17, 2023 11:12:28 GMT+08:00	Permissions More
	EXTERNAL		ORC	0 B	---	Jan 17, 2023 11:13:27 GMT+08:00	Jan 17, 2023 11:13:27 GMT+08:00	Permissions More
	EXTERNAL		ORC	0 B	---	Mar 10, 2023 10:44:22 GMT+08:00	Mar 10, 2023 10:44:22 GMT+08:00	Permissions More

- To import data on the **SQL Editor** page:
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the navigation tree on the left of **SQL Editor**, click **Databases** to see all databases. Click the database where the target table belongs. The table list is displayed.
 - c. Click  on the right of the table and choose **Import** from the shortcut menu. The **Import** page is displayed.

Figure 5-12 SQL editor - importing data



Step 2 In the **Import** dialog box, set the parameters based on [Table 5-10](#).

Table 5-10 Description

Parameter	Description	Example
Databases	Database where the current table is located.	-
Table Name	Name of the current table.	-
Queues	Queue where the imported data will be used	-
File Format	Format of the data source file to be imported. The CSV, Parquet, ORC, JSON, and Avro formats are supported. Encoding format. Only UTF-8 is supported.	CSV
Path	<p>You can directly enter a path or click  and select an OBS path. If no bucket is available, you can directly switch to the OBS management console and create an OBS bucket.</p> <ul style="list-style-type: none"> • When creating an OBS table, you must specify a folder as the directory. If a file is specified, data import may be failed. • If a folder and a file have the same name in the OBS directory, the file path is preferred as the path of the data to be imported. <p>NOTE The path can be a file or folder.</p>	obs://DLI/ sampledat a.csv

Parameter	Description	Example
Table Header: No/Yes	<p>This parameter is valid only when File Format is set to CSV. Whether the data source to be imported contains the table header.</p> <p>Click Advanced Settings and select the check box next to Table Header: No. If the check box is selected, the table header is displayed. If the check box is deselected, no table header is displayed.</p>	-
User-defined Delimiter	<p>This parameter is valid only when File Format is set to CSV and you select User-defined Delimiter.</p> <p>The following delimiters are supported:</p> <ul style="list-style-type: none">• Comma (,)• Vertical bar ()• Tab character (\t)• Others: Enter a user-defined delimiter.	Default value: comma (,)
User-defined Quotation Character	<p>This parameter is valid only when File Format is set to CSV and User-defined Quotation Character is selected.</p> <p>The following quotation characters are supported:</p> <ul style="list-style-type: none">• Single quotation mark (')• Double quotation marks (")• Others: Enter a user-defined quotation character.	Default value: double quotation marks (")
User-defined Escape Character	<p>This parameter is valid only when File Format is set to CSV and you select User-defined Escape Character.</p> <p>The following escape characters are supported:</p> <ul style="list-style-type: none">• Backslash (\)• Others: Enter a user-defined escape character.	Default value: backslash (\)
Date Format	<p>This parameter is valid only when File Format is set to CSV or JSON.</p> <p>This parameter specifies the format of the date in the table and is valid only Advanced Settings is selected. The default value is yyyy-MM-dd. For definition of characters involved in the date pattern, see Table 3 in Importing Data to the Table.</p>	2000-01-0 1

Parameter	Description	Example
Timestamp Format	This parameter is valid only when File Format is set to CSV or JSON . This parameter specifies the format of the timestamp in the table and is valid only Advanced Settings is selected. The default value is yyyy-MM-dd HH:mm:ss . For definition of characters involved in the time pattern, see Table 3 in Importing Data to the Table .	2000-01-01 09:00:00
Error Records Path	This parameter is valid only when File Format is set to CSV or JSON . The parameter specifies the error data is stored in the corresponding OBS path and is valid only Advanced Settings is selected.	obs://DLI/

Step 3 Click **OK**.

Step 4 You can view the imported data in either of the following ways:

 **NOTE**

Currently, only the first 10 records are displayed.

- Choose **Data Management > Databases and Tables** in the navigation pane of the console. Locate the row that contains the database where the target table belongs and click **More > View Properties** in the **Operation** column. In the displayed dialog box, click the **Preview** tab to view the imported data.
- In the **Databases** tab of the **SQL Editor**, click the database name to go to the table list. Click  on the right of a table name and choose **View Properties** from the shortcut menu. In the displayed dialog box, click **Preview** to view the imported data.

Step 5 (Optional) View the status and execution result of the importing job on the **Job Management > SQL Jobs** page.

----End

5.1.8 Exporting Data from DLI to OBS

You can export data from a DLI table to OBS. During the export, a folder is created in OBS or the content in the existing folder is overwritten.

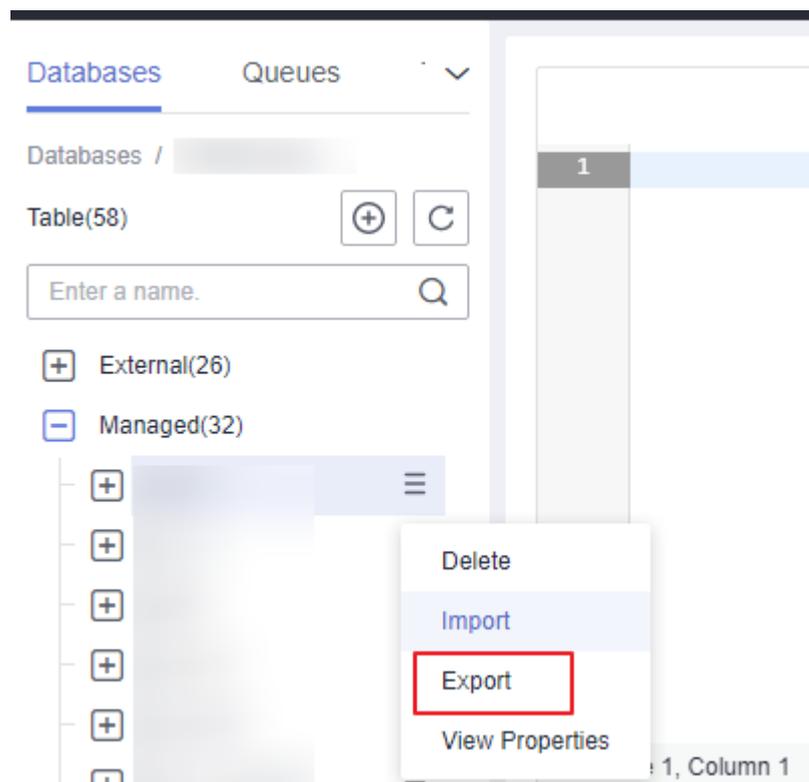
Precautions

- The exported file can be in JSON format, and the text format can only be UTF-8.
- Only the data in the DLI table (the table type is **Managed**) can be exported to the OBS bucket, and the export path must be specified to the folder level.
- Data can be exported across accounts. That is, after account B authorizes account A, account A can export data to the OBS path of account B if account A has the permission to read the metadata and permission information about the OBS bucket of account B and read and write the path.

Procedure

- Step 1** You can export data on either the **Data Management** page or the **SQL Editor** page.
- To export data on the **Data Management** page:
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. Click the name of the database corresponding to the table where data is to be exported to switch to the **Manage Tables** page.
 - c. Select the corresponding table (DLI table) and choose **More > Export** in the **Operation** column. The **Export Data** page is displayed.
 - To export data on the **SQL Editor** page:
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the navigation tree on the left, click **Databases** to see all databases. Click the database name corresponding to the table to which data is to be exported. The tables are displayed.
 - c. Click  on the right of the managed table (DLI table) whose data is to be exported, and choose **Export** from the shortcut menu.

Figure 5-13 Exporting the table



- Step 2** In the displayed **Export Data** dialog box, specify parameters by referring to [Table 5-11](#).

Figure 5-14 Exporting data

Export Result
×

★ Databases db1

★ Table Name src

★ Data Format json

Queues

Compression Format

★ Storage Path

★ Export Mode
 New OBS directory
 Existing OBS directory (Overwritten)

Table Header
 No
 Yes

Table 5-11 Parameter description

Parameter	Description
Databases	Database where the current table is located.
Table Name	Name of the current table.
Data Format	Format of the file storing data to be exported. Formats other than JSON will be supported in later versions.
Queue	Select a queue.
Compression Format	Compression format of the data to be exported. The following compression formats are supported: <ul style="list-style-type: none"> • none • bzip2 • deflate • gzip
Storage Path	<ul style="list-style-type: none"> • Enter or select an OBS path. • The export path must be a folder that does not exist in the OBS bucket. Specifically, you need to create a folder in the target OBS directory. • The folder name cannot contain the special characters of \/:*?"<> , and cannot start or end with a dot (.).

Parameter	Description
Export Mode	Storage mode of the data to be exported. <ul style="list-style-type: none">• New OBS directory: If the specified export directory exists, an error is reported and the export operation cannot be performed.• Existing OBS directory (Overwritten): If you create a file in the specified directory, the existing file will be overwritten.
Table Header: No/Yes	Whether the data to be exported contains the table header.

Step 3 Click **OK**.

Step 4 (Optional) You can view the job status (indicated by **Status**), statements (indicated by **Statement**), and other information about exporting jobs on the **SQL Jobs** page.

1. Select **EXPORT** from the **Job Type** drop-down list box and specify the time range for exporting data. The jobs meeting the requirements are displayed in the job list.
2. Click  to view details about an exporting job.

----End

5.1.9 Viewing Metadata

Metadata Description

- Metadata is used to define data types. It describes information about the data, including the source, size, format, and other data features. In database fields, metadata interprets data content in the data warehouse.
- When you create a table, metadata is defined, consisting of the column name, type, and description.
- The **Metadata** page displays information about the target table, including **Column Name**, **Column Type**, **Data Type**, and **Description**.

Procedure

You can view metadata on either the **Data Management** page or the **SQL Editor** page.

- To view metadata on the **Data Management** page:
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. On the displayed **Data Management** page, click the name of the database where the target table whose data you want to export resides to switch to the **Manage Tables** page.
 - c. Click **More** in the **Operation** column of the target table and select **View Properties**. In the **Metadata** tab, view the metadata of the table.

- To view metadata on the **SQL Editor** page:
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the navigation pane of the displayed **SQL Editor** page, click **Databases**.
 - c. Click the corresponding database name to view the tables in the database.
 - d. Click  on the right of the table and choose **View Properties** from the shortcut menu. On the **Metadata** tab page, view the metadata of the table.

5.1.10 Previewing Data

The **Preview** page displays the first 10 records in the table.

Procedure

You can preview data on either the **Data Management** page or the **SQL Editor** page.

- To preview data on the **Data Management** page:
 - a. On the left of the management console, choose **Data Management > Databases and Tables**.
 - b. On the displayed **Data Management** page, click the name of the database where the target table whose data you want to export resides to switch to the **Manage Tables** page.
 - c. Click **More** in the **Operation** column of the target table and select **View Properties**.
 - d. Click the **Preview** tab to preview the table data.
- To preview data on the **SQL Editor** page:
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the navigation pane of the displayed **SQL Editor** page, click **Databases**.
 - c. Click the corresponding database name to view the tables in the database.
 - d. Click  on the right of the corresponding table, choose **View Properties** from the list menu, and click the **Preview** tab to preview the data of the table.

5.1.11 Managing Tags

Tag Management

A tag is a key-value pair that you can customize to identify cloud resources. It helps you to classify and search for cloud resources. A tag consists of a tag key and a tag value. If you use tags in other cloud services, you are advised to create the same tag (key-value pairs) for cloud resources used by the same business to keep consistency.

DLI supports the following two types of tags:

- Resource tags: non-global tags created on DLI
- Predefined tags: global tags created on Tag Management Service (TMS).
For more information about predefined tags, see [Tag Management Service User Guide](#).

This section describes how to add, modify, and delete tags for databases and tables.

Database Tags

- Step 1** In the navigation pane on the left, choose **Data Management > Databases and Tables**.
- Step 2** Locate the row that contains the target database, and click **More > Tags** in the **Operation** column.
- Step 3** The tag management page is displayed, and the tags (if there are) are displayed.
- Step 4** On the displayed page, click **Add/Edit Tag**. The **Add/Edit Tag** dialog box is displayed.

Enter a tag key and a tag value in the text boxes and click **Add**.

Figure 5-15 Adding/Editing tags

Add/Edit Tag ×

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#) ↻

To add a tag, enter a tag key and a tag value below.

test = test ×

dli 001 Add

19 tags available for addition.

OK Cancel

Table 5-12 Tag parameters

Parameter	Description
Tag key	<p>You can specify the tag key in either of the following ways:</p> <ul style="list-style-type: none">Click the text box for tag key and select a predefined tag key from the drop-down list. To add a predefined tag, you need to create one on TMS and then select it from the Tag key drop-down list. You can click View predefined tags to go to the Predefined Tags page of the TMS console. Then, click Create Tag in the upper corner of the page to create a predefined tag. For details, see Creating Predefined Tags in <i>Tag Management Service User Guide</i>.Enter a tag key in the text box. <p>NOTE A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed, but the value cannot start or end with a space or start with _sys_.</p>
Tag value	<p>You can specify the tag value in either of the following ways:</p> <ul style="list-style-type: none">Click the tag value text box and select a predefined tag value from the drop-down list.Enter a tag value in the text box. <p>NOTE A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed. The value cannot start or end with a space.</p>

NOTE

- A maximum of 20 tags can be added.
- Only one tag value can be added to a tag key.
- The key name in each resource must be unique.

Step 5 Click **OK**. The database tag is added.

To delete a tag, click **Delete** in the **Operation** column of the target tag.

----End

Table Tags

Step 1 In the navigation pane on the left, choose **Data Management > Databases and Tables**.

Step 2 Click a database name to view the tables in the database.

Step 3 Locate the row that contains the target table and click **More > Tag** in the **Operation** column.

Step 4 The tag management page is displayed, and the tags (if there are) are displayed.

Step 5 On the displayed page, click **Add/Edit Tag**. The **Add/Edit Tag** dialog box is displayed.

Enter a tag key and a tag value in the text boxes and click **Add**

Figure 5-16 Adding/Editing tags

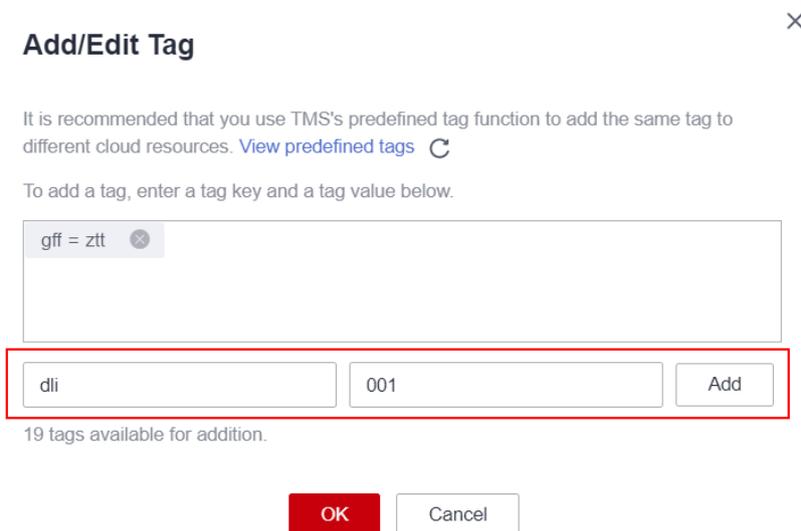


Table 5-13 Tag parameters

Parameter	Description
Tag key	<p>You can specify the tag key in either of the following ways:</p> <ul style="list-style-type: none"> Click the text box for tag key and select a predefined tag key from the drop-down list. To add a predefined tag, you need to create one on TMS and then select it from the Tag key drop-down list. You can click View predefined tags to go to the Predefined Tags page of the TMS console. Then, click Create Tag in the upper corner of the page to create a predefined tag. For details, see Creating Predefined Tags in <i>Tag Management Service User Guide</i>. Enter a tag key in the text box. <p>NOTE A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_ :=+ - @) are allowed, but the value cannot start or end with a space or start with _sys_.</p>

Parameter	Description
Tag value	<p>You can specify the tag value in either of the following ways:</p> <ul style="list-style-type: none">• Click the tag value text box and select a predefined tag value from the drop-down list.• Enter a tag value in the text box. <p>NOTE A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:+=-@) are allowed. The value cannot start or end with a space.</p>

 **NOTE**

- A maximum of 20 tags can be added.
- Only one tag value can be added to a tag key.
- The key name in each resource must be unique.

Step 6 Click **OK**. The table tag is added.

To delete a tag, click **Delete** in the **Operation** column of the target tag.

----End

5.2 Package Management

5.2.1 Overview

Package management provides the following functions:

- [Managing Package Permissions](#)
- [Creating a Package](#)
- [Deleting a Package](#)

 **NOTE**

You can delete program packages in batches.

- [Modifying the Owner](#)

Constraints

- A package can be deleted, but a package group cannot be deleted.
- The following types of packages can be uploaded:
 - **JAR**: JAR file
 - **PyFile**: User Python file
 - **File**: User file
 - **ModelFile**: User AI model file

Package Management Page

Table 5-14 Package management parameters

Parameter	Description
Group Name	Name of the group to which the package belongs. If the package is not grouped, -- is displayed.
Package Name	Name of a package.
Owner	Name of the user who uploads the package.
Type	Type of a package. The following package types are supported: <ul style="list-style-type: none">• JAR: JAR file• PyFile: User Python file• File: User file• ModelFile: User AI model file
Status	Status of the package to be created. <ul style="list-style-type: none">• Uploading: The file is being uploaded.• Finished: The resource package has been uploaded.• Failed: The resource package upload failed.
Created	Time when a package is created.
Updated	Time when the package is updated.
Operation	Manage Permissions: Manage user permissions for a package. Delete: Delete the package. More: <ul style="list-style-type: none">• Modify Owner: Modify the owner of the package.• Tags: Add or edit package tags.

5.2.2 Managing Permissions on Packages and Package Groups

Scenario

- You can isolate package groups or packages allocated to different users by setting permissions to ensure data query performance.
- The administrator and the owner of a package group or package have all permissions. You do not need to set permissions and the permissions cannot be modified by other users.
- When you set permissions on a package group or a package to a new user, the region of the user group the user belong to must have the Tenant Guest permission. For details about the Tenant Guest permission and how to apply for the permission, see [Creating a User Group and Assigning Permissions](#) and [System Permissions](#).

On the **Package Management** page, click **Manage Permissions** in the **Operation** column of the target package. On the displayed **User Permission Info** page, you can grant permissions for the package group or package, set and revoke user permissions.

 **NOTE**

- If you select a group when creating a package, you can manage permissions of the corresponding program package group.
- If you select **No grouping** when creating a package, you can manage permissions of the corresponding package.

Granting Permissions on Package Groups/Packages

Click **Grant Permission** in the upper right corner of the page.

- Granting permissions on package groups

Figure 5-17 Granting permissions on package groups

Grant Permission

* Username

Select the permissions to be granted to the user

Select all
 Use Group
 Update Group
 Get Group

Delete Group
 Grant Permission
 Revoke Permission

View Other User's Permissions

Table 5-15 Permission parameters

Parameter	Description
Username	Name of the authorized IAM user. NOTE The username is the name of an existing IAM user.

Parameter	Description
Select the permissions to be granted to the user	<ul style="list-style-type: none"> • Use Group: This permission allows you to use the package of this group. • Update Group: This permission allows you to update the packages in the group, including creating a package in the group. • Query Group: This permission allows you to query the details of a package in a group. • Delete Group: This permission allows you to delete the package of the group. • Grant Permission: This permission allows you to grant group permissions to other users. • Revoke Permission: This permission allows you to revoke the group permissions that other users have but cannot revoke the group owner's permissions. • View Other User's Permissions: This permission allows you to view the group permissions of other users.

- Granting permissions on packages

Figure 5-18 Granting permissions on package groups

Grant Permission

* Username

Select the permissions to be granted to the user

Select all
 Use Resource
 Update Resource
 Get Resource

Delete Resource
 Grant Permission
 Revoke Permission

View Other User's Permissions

Table 5-16 Permission parameters

Parameter	Description
Username	Name of the authorized IAM user. NOTE The username is the name of an existing IAM user.

Parameter	Description
Select the permissions to be granted to the user	<ul style="list-style-type: none">● Use Package: This permission allows you to use the package.● Update Package: This permission allows you to update the package.● Query Package: This permission allows you to query the package.● Delete Package: This permission allows you to delete the package.● Grant Permission: This permission allows you to grant package permissions to other users.● Revoke Permission: This permission allows you to revoke the package permissions that other users have but cannot revoke the package owner's permissions.● View Other User's Permissions: This permission allows you to view the package permissions of other users.

Setting Permissions on Package Groups and Packages

Click **Set Permission** in the **Operation** column of the sub-user to modify the permission of the user. [Table 5-15](#) and [Table 5-16](#) list the detailed permission descriptions.

If the **Set Permission** button is gray, you do not have the permission to modify the package group or package. You can apply to the administrator, group owner, or other users who have the permissions on granting and revoking permissions of package groups or packages.

Revoking Permissions on Package Groups and Packages

Click **Revoke Permission** in the **Operation** column of a sub-user to revoke the user's permissions. After the operation, the sub-user does not have any permission on the package group or package.

Permissions Description

- Package group permissions
Querying permissions. A group owner can view the created package group and all packages in the group, and can also view package groups on which they have all permissions.
A package group is a unit. If you select a group when creating a package, you can grant only the permissions of the package group to other users.
- Package permissions
Querying permissions. A package owner can view the created packages, and can also view packages on which they have all permissions.

5.2.3 Creating a Package

DLI allows you to submit program packages in batches to the general-use queue for running.

NOTE

If you need to update a package, you can use the same package or file to upload it to the same location (in the same group) on DLI to overwrite the original package or file.

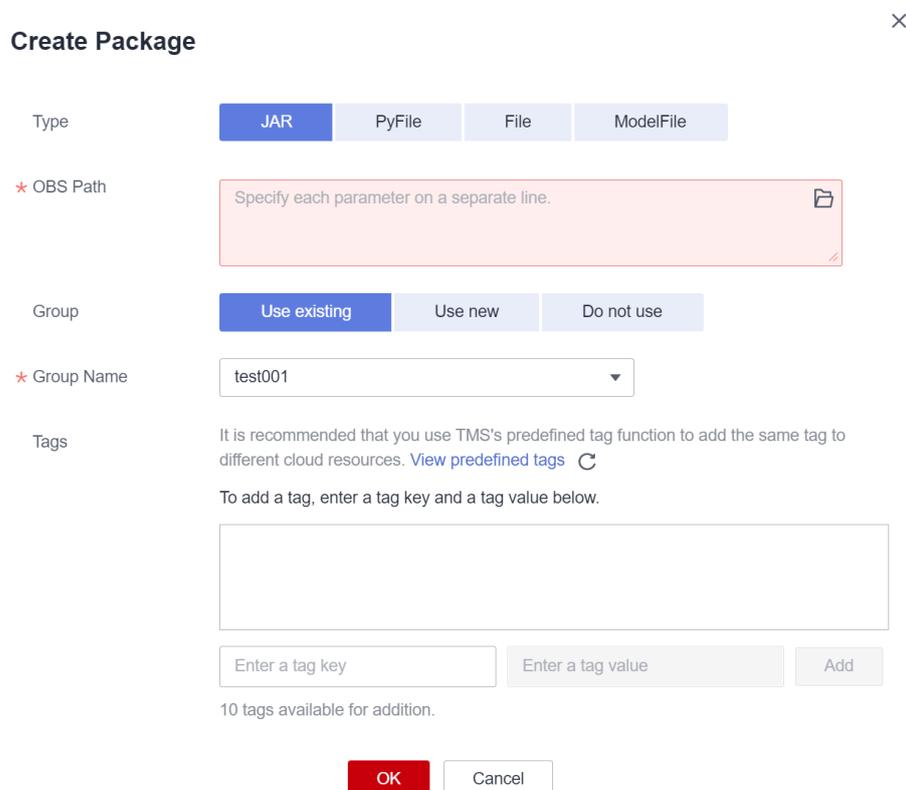
Prerequisites

All software packages must be uploaded to OBS for storage in advance.

Procedure

1. On the left of the management console, choose **Data Management > Package Management**.
2. On the **Package Management** page, click **Create** in the upper right corner to create a package.
3. In the displayed **Create Package** dialog box, set related parameters by referring to [Table 5-17](#).

Figure 5-19 Creating a package



Create Package ×

Type JAR PyFile File ModelFile

* OBS Path 

Group Use existing Use new Do not use

* Group Name

Tags It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#) 
To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

Table 5-17 Parameter description

Parameter	Description
Package Type	The following package types are supported: <ul style="list-style-type: none">● JAR: JAR file● PyFile: User Python file● File: User file● ModelFile: User AI model file
Package File Path	Select the OBS path of the corresponding packages. NOTE <ul style="list-style-type: none">● The packages must be uploaded to OBS for storage in advance.● Only files can be selected.
Group Policy	You can select Use existing group , Use new group , or No grouping .
Group Name	<ul style="list-style-type: none">● Use existing group: Select an existing group.● Use new group: Enter a custom group name.● No grouping: No need to select or enter a group name. NOTE <ul style="list-style-type: none">● If you select a group, the permission management refers to the permissions of the corresponding package group.● If no group is selected, the permission management refers to the permissions of the corresponding package. For details about how to manage permissions on package groups and packages, see Managing Permissions on Packages and Package Groups .
Tag	Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS). For details, see Tag Management Service User Guide . NOTE <ul style="list-style-type: none">● A maximum of 20 tags can be added.● Only one tag value can be added to a tag key.● The key name in each resource must be unique. <ul style="list-style-type: none">● Tag key: Enter a tag key name in the text box. NOTE A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:+=@) are allowed, but the value cannot start or end with a space or start with <code>_sys_</code>.● Tag value: Enter a tag value in the text box. NOTE A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:+=@) are allowed. The value cannot start or end with a space.

4. Click **OK**.

After a package is created, you can view and select the package for use on the **Package Management** page.

5.2.4 Deleting a Package

You can delete a package based on actual conditions.

Procedure

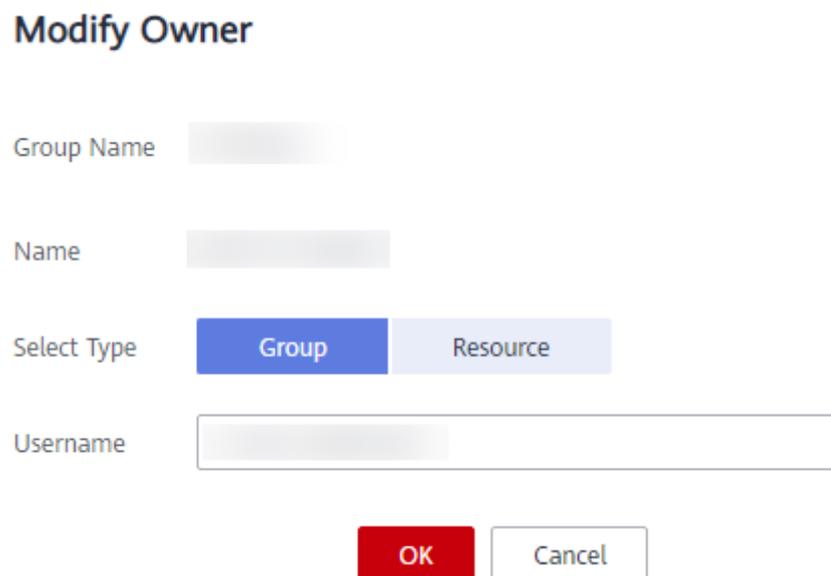
1. On the left of the management console, choose **Data Management > Package Management**.
2. Click **Delete** in the **Operation** column of the package to be deleted.
3. In the dialog box that is displayed, click **Yes**.

5.2.5 Modifying the Owner

To change the owner of a package, click **More > Modify Owner** in the **Operation** column of a package on the **Package Management** page.

- If the package has been grouped, you can modify the owner of the **Group** or **Resource** of it.

Figure 5-20 Modifying the package owner



Modify Owner

Group Name

Name

Select Type Group Resource

Username

- If the package has not been grouped, change its owner directly.

Figure 5-21 Modifying the owner of a package

Table 5-18 Description

Parameter	Description
Group Name	<ul style="list-style-type: none"> If you select a group when creating a package, the name of the group is displayed. If no group is selected when creating a package, this parameter is not displayed.
Name	Name of a package.
Select Type	<ul style="list-style-type: none"> If you select a group when creating a package, you can change the owner of the group or package. If no group is selected when creating a package, this parameter is not displayed.
Username	Name of the package owner. NOTE The username is the name of an existing IAM user.

5.2.6 Built-in Dependencies

DLI built-in dependencies are provided by the platform by default. In case of conflicts, you do not need to upload them when packaging JAR packages of Spark or Flink Jar jobs.

Spark 2.4.5 Dependencies

- JavaEWAH-1.1.7.jar
- RoaringBitmap-0.7.45.jar
- ST4-4.3.1.jar
- accessors-smart-1.2.jar
- activation-1.1.1.jar
- aircompressor-0.16.jar

- alluxio-2.3.1-luxor-SNAPSHOT-client.jar
- annotations-17.0.0.jar
- antlr-2.7.7.jar
- antlr-runtime-3.4.jar
- antlr4-runtime-4.8-1.jar
- aopalliance-1.0.jar
- aopalliance-repackaged-2.4.0-b34.jar
- apache-log4j-extras-1.2.17.jar
- arpack_combined_all-0.1.jar
- arrow-format-0.12.0.jar
- arrow-memory-0.12.0.jar
- arrow-vector-0.12.0.jar
- asm-5.0.4.jar
- audience-annotations-0.5.0.jar
- automaton-1.11-8.jar
- avro-1.8.2.jar
- avro-ipc-1.8.2.jar
- avro-mapred-1.8.2.jar
- java-sdk-bundle-1.11.856.jar
- base64-2.3.8.jar
- bcpkix-jdk15on-1.66.jar
- bcprov-jdk15on-1.67.jar
- bonecp-0.8.0.RELEASE.jar
- breeze-macros_2.11-0.13.2.jar
- breeze_2.11-0.13.2.jar
- calcite-avatica-1.2.0-incubating.jar
- chill-java-0.9.3.jar
- chill_2.11-0.9.3.jar
- commons-beanutils-1.9.4.jar
- commons-cli-1.2.jar
- commons-codec-1.15.jar
- commons-collections-3.2.2.jar
- commons-collections4-4.2.jar
- commons-compiler-3.0.9.jar
- commons-compress-1.4.1.jar
- commons-configuration2-2.1.1.jar
- commons-crypto-1.0.0-20191105.jar
- commons-daemon-1.0.13.jar
- commons-dbcp2-2.7.0.jar
- commons-httpclient-3.1.jar

- commons-io-2.5.jar
- commons-lang-2.6.jar
- commons-lang3-3.5.jar
- commons-logging-1.2.jar
- commons-math3-3.4.1.jar
- commons-net-3.1.jar
- commons-pool2-2.8.0.jar
- commons-text-1.3.jar
- compress-lzf-1.0.3.jar
- core-1.1.2.jar
- crypter-0.0.6.jar
- curator-client-4.2.0.jar
- curator-framework-4.2.0.jar
- curator-recipes-2.7.1.jar
- datanucleus-api-jdo-3.2.6.jar
- datanucleus-core-3.2.10.jar
- datanucleus-rdbms-3.2.9.jar
- derby-10.14.2.0.jar
- dnsjava-2.1.7.jar
- ecj-3.21.0.jar
- ehcache-3.3.1.jar
- expiringmap-0.5.9.jar
- fastutil-8.2.3.jar
- flatbuffers-java-1.9.0.jar
- fst-2.50.jar
- generex-1.0.2.jar
- geronimo-jcache_1.0_spec-1.0-alpha-1.jar
- gson-2.2.4.jar
- guava-14.0.1.jar
- guice-4.0.jar
- guice-servlet-4.0.jar
- hadoop-annotations-3.1.1-ei-302002.jar
- hadoop-auth-3.1.1-ei-302002.jar
- hadoop-3.1.1-ei-302002.jar
- hadoop-client-3.1.1-ei-302002.jar
- hadoop-common-3.1.1-ei-302002.jar
- hadoop-hdfs-3.1.1-ei-302002.jar
- hadoop-hdfs-client-3.1.1-ei-302002.jar
- hadoop-3.1.1-46.jar
- hadoop-mapreduce-client-common-3.1.1-ei-302002.jar

- hadoop-mapreduce-client-core-3.1.1-ei-302002.jar
- hadoop-mapreduce-client-jobclient-3.1.1-ei-302002.jar
- hadoop-minikdc-3.1.1-ei-302002.jar
- hadoop-yarn-api-3.1.1-ei-302002.jar
- hadoop-yarn-client-3.1.1-ei-302002.jar
- hadoop-yarn-common-3.1.1-ei-302002.jar
- hadoop-yarn-registry-3.1.1-ei-302002.jar
- hadoop-yarn-server-applicationhistoryservice-3.1.1-ei-302002.jar
- hadoop-yarn-server-common-3.1.1-ei-302002.jar
- hadoop-yarn-server-resourcemanager-3.1.1-ei-302002.jar
- hadoop-yarn-server-web-proxy-3.1.1-ei-302002.jar
- hamcrest-core-1.3.jar
- hive-common-1.2.1-2.0.0.dli-20220528.090500-402.jar
- hive-exec-1.2.1-2.0.0.dli-20220528.090521-401.jar
- hive-metastore-1.2.1-2.0.0.dli-20220528.090509-402.jar
- hive-shims-0.23-1.2.1-2.0.0.dli-20220528.090445-403.jar
- hive-shims-1.2.1-2.0.0.dli-20220528.090455-403.jar
- hive-shims-common-1.2.1-2.0.0.dli-20220528.090441-404.jar
- hive-shims-scheduler-1.2.1-2.0.0.dli-20220528.090450-403.jar
- hk2-api-2.4.0-b34.jar
- hk2-locator-2.4.0-b34.jar
- hk2-utils-2.4.0-b34.jar
- hppc-0.7.2.jar
- htrace-core4-4.2.0-incubating-1.0.0.jar
- httpclient-4.5.6.jar
- httpcore-4.4.10.jar
- ivy-2.4.0.jar
- jackson-annotations-2.11.4.jar
- jackson-core-2.11.4.jar
- jackson-core-asl-1.9.13-atlassian-4.jar
- jackson-databind-2.11.4.jar
- jackson-dataformat-yaml-2.11.4.jar
- jackson-datatype-jsr310-2.11.2.jar
- jackson-jaxrs-base-2.10.3.jar
- jackson-jaxrs-json-provider-2.10.3.jar
- jackson-mapper-asl-1.9.13-atlassian-4.jar
- jackson-module-jaxb-annotations-2.10.3.jar
- jackson-module-paranamer-2.11.4.jar
- jackson-module-scala_2.11-2.11.4.jar
- jakarta.activation-api-1.2.1.jar

- jakarta.xml.bind-api-2.3.2.jar
- janino-3.0.9.jar
- java-util-1.9.0.jar
- java-xmlbuilder-1.1.jar
- javassist-3.18.1-GA.jar
- javax.annotation-api-1.2.jar
- javax.inject-1.jar
- javax.inject-2.4.0-b34.jar
- javax.servlet-api-3.1.0.jar
- javax.ws.rs-api-2.0.1.jar
- javolution-5.3.1.jar
- jaxb-api-2.2.11.jar
- jcip-annotations-1.0-1.jar
- jcl-over-slf4j-1.7.30.jar
- jdo-api-3.0.1.jar
- jersey-client-2.23.1.jar
- jersey-common-2.23.1.jar
- jersey-container-servlet-2.23.1.jar
- jersey-container-servlet-core-2.23.1.jar
- jersey-guava-2.23.1.jar
- jersey-media-jaxb-2.23.1.jar
- jersey-server-2.23.1.jar
- jets3t-0.9.4.jar
- jettison-1.1.jar
- jetty-http-9.4.34.v20201102.jar
- jetty-io-9.4.34.v20201102.jar
- jetty-security-9.4.34.v20201102.jar
- jetty-server-9.4.34.v20201102.jar
- jetty-servlet-9.4.34.v20201102.jar
- jetty-util-9.4.34.v20201102.jar
- jetty-util-ajax-9.4.34.v20201102.jar
- jetty-webapp-9.4.34.v20201102.jar
- jetty-xml-9.4.34.v20201102.jar
- joda-time-2.9.3.jar
- jodd-core-3.5.2.jar
- json-20200518.jar
- json-io-2.5.1.jar
- json-sanitizer-1.2.1.jar
- json-smart-2.3.jar
- json4s-ast_2.11-3.5.3.jar

- json4s-core_2.11-3.5.3.jar
- json4s-jackson_2.11-3.5.3.jar
- json4s-scalap_2.11-3.5.3.jar
- jsp-api-2.1.jar
- jsr305-1.3.9.jar
- jta-1.1.jar
- jtransforms-2.4.0.jar
- jts-core-1.16.1.jar
- jul-to-slf4j-1.7.30.jar
- junit-4.11.jar
- kerb-admin-1.0.1.jar
- kerb-client-1.0.1.jar
- kerb-common-1.0.1.jar
- kerb-core-1.0.1.jar
- kerb-crypto-1.0.1.jar
- kerb-identity-1.0.1.jar
- kerb-server-1.0.1.jar
- kerb-simplekdc-1.0.1.jar
- kerb-util-1.0.1.jar
- kerby-asn1-1.0.1.jar
- kerby-config-1.0.1.jar
- kerby-pkix-1.0.1.jar
- kerby-util-1.0.1.jar
- kerby-xdr-1.0.1.jar
- kryo-shaded-4.0.2.jar
- kubernetes-client-5.4.1-20211025.jar
- kubernetes-model-admissionregistration-5.4.1-20211025.jar
- kubernetes-model-apiextensions-5.4.1-20211025.jar
- kubernetes-model-apps-5.4.1-20211025.jar
- kubernetes-model-autoscaling-5.4.1-20211025.jar
- kubernetes-model-batch-5.4.1-20211025.jar
- kubernetes-model-certificates-5.4.1-20211025.jar
- kubernetes-model-common-5.4.1-20211025.jar
- kubernetes-model-coordination-5.4.1-20211025.jar
- kubernetes-model-core-5.4.1-20211025.jar
- kubernetes-model-discovery-5.4.1-20211025.jar
- kubernetes-model-events-5.4.1-20211025.jar
- kubernetes-model-extensions-5.4.1-20211025.jar
- kubernetes-model-flowcontrol-5.4.1-20211025.jar
- kubernetes-model-metrics-5.4.1-20211025.jar

- kubernetes-model-networking-5.4.1-20211025.jar
- kubernetes-model-node-5.4.1-20211025.jar
- kubernetes-model-policy-5.4.1-20211025.jar
- kubernetes-model-rbac-5.4.1-20211025.jar
- kubernetes-model-scheduling-5.4.1-20211025.jar
- kubernetes-model-storageclass-5.4.1-20211025.jar
- leveldbjni-all-1.8-20191105.jar
- libfb303-0.9.3.jar
- libthrift-0.12.0.jar
- log4j-1.2.17-cloudera1.jar
- log4j-rolling-appender-20131024-2017.jar
- logging-interceptor-3.14.9.jar
- lucene-analyzers-common-7.7.2.jar
- lucene-core-7.7.2.jar
- lucene-queryparser-7.7.2.jar
- lucene-sandbox-7.7.2.jar
- luxor-encrypt-2.0.0-20220623.010726-213.jar
- luxor-fs3-2.0.0-20220623.010750-209.jar
- luxor-obs-fs3-2.0.0-20220623.010756-209.jar
- luxor-rpc_2.11-2.0.0-20220623.010737-182.jar
- luxor-transport-2.0.0-20220623.010744-71.jar
- lz4-java-1.7.1.jar
- machinist_2.11-0.6.1.jar
- macro-compat_2.11-1.1.1.jar
- metrics-core-3.1.5.jar
- metrics-graphite-3.1.5.jar
- metrics-jmx-4.1.12.1.jar
- metrics-json-3.1.5.jar
- metrics-jvm-3.1.5.jar
- minlog-1.3.0.jar
- mssql-jdbc-6.2.1.jre7.jar
- netty-all-4.1.51.Final.jar
- nimbus-jose-jwt-8.19.jar
- objenesis-2.5.1.jar
- okhttp-3.14.9.jar
- okio-1.17.2.jar
- opencsv-2.3.jar
- opencsv-4.6.jar
- opencv-4.3.0-2.jar
- orc-core-1.6.8-nohive.jar

- orc-mapreduce-1.6.8-nohive.jar
- orc-shims-1.6.8.jar
- oro-2.0.8.jar
- osgi-resource-locator-1.0.1.jar
- paranamer-2.8.jar
- parquet-column-1.12.2.jar
- parquet-common-1.12.2.jar
- parquet-encoding-1.12.2.jar
- parquet-format-structures-1.12.2.jar
- parquet-hadoop-1.12.2.jar
- parquet-hadoop-bundle-1.6.0.jar
- parquet-jackson-1.12.2.jar
- postgresql-42.2.14.jar
- protobuf-java-2.5.0.jar
- py4j-0.10.7.jar
- pyrolite-4.13.jar
- re2j-1.1.jar
- scala-compiler-2.11.12.jar
- scala-library-2.11.12.jar
- scala-parser-combinators_2.11-1.1.2.jar
- scala-reflect-2.11.12.jar
- scala-xml_2.11-1.0.5.jar
- secComponentApi-1.0.6.jar
- shapeless_2.11-2.3.2.jar
- shims-0.7.45.jar
- slf4j-api-1.7.30.jar
- slf4j-log4j12-1.7.30.jar
- snakeyaml-1.26.jar
- snappy-java-1.1.8.2.jar
- solr-core-7.7.2.jar
- solr-solrj-7.7.2.jar
- spark-avro_2.11-2.4.5.0100-2.0.0.dli-20220617.085536-9.jar
- spark-avro_2.11-4.0.0.jar
- spark-catalyst_2.11-2.4.5.0100-2.0.0.dli-20220617.085405-16.jar
- spark-core_2.11-2.4.5.0100-2.0.0.dli-20220617.085327-16.jar
- spark-graphx_2.11-2.4.5.0100-.0.dli-20220617.085336-16.jar
- spark-hive_2.11-2.4.5.0100-2.0.0.dli-20220617.085423-16.jar
- spark-kubernetes_2.11-2.4.5.0100-2.0.0.dli-20220617.085519-16.jar
- spark-kvstore_2.11-2.4.5.0100-2.0.0.dli-20220617.085249-16.jar
- spark-launcher_2.11-2.4.5.0100-2.0.0.dli-20220617.085435-16.jar

- spark-mllib-local_2.11-2.4.5.0100-2.0.0.dli-20220617.085349-16.jar
- spark-mllib_2.11-2.4.5.0100-2.0.0.dli-20220617.085342-16.jar
- spark-network-common_2.11-2.4.5.0100-2.0.0.dli-20220617.085254-16.jar
- spark-network-shuffle_2.11-2.4.5.0100-2.0.0.dli-20220617.085300-16.jar
- spark-om_2.11-2.4.5.0100-2.0.0.dli-20220617.085316-16.jar
- spark-repl_2.11-2.4.5.0100-2.0.0.dli-20220617.085430-16.jar
- spark-sketch_2.11-2.4.5.0100-2.0.0.dli-20220617.085243-16.jar
- spark-sql_2.11-2.4.5.0100-2.0.0.dli-20220617.085414-16.jar
- spark-streaming_2.11-2.4.5.0100-.0.dli-20220617.085359-16.jar
- spark-tags_2.11-2.4.5.0100-2.0.0.dli-20220617.085322-16.jar
- spark-unsafe_2.11-2.4.5.0100-2.0.0.dli-20220617.085311-16.jar
- spark-uquery_2.11-2.4.5.0100-2.0.0.dli-SNAPSHOT.jar
- spark-yarn_2.11-2.4.5.0100-2.0.0.dli-20220617.085531-16.jar
- spire-macros_2.11-0.13.0.jar
- spire_2.11-0.13.0.jar
- stax-api-1.0-2.jar
- stax2-api-3.1.4.jar
- stream-2.7.0.jar
- stringtemplate-3.2.1.jar
- threeten-extra-1.5.0.jar
- tink-1.6.0.jar
- token-provider-1.0.1.jar
- tomcat-api-9.0.39.jar
- tomcat-el-api-9.0.39.jar
- tomcat-jasper-9.0.39.jar
- tomcat-jasper-el-9.0.39.jar
- tomcat-jsp-api-9.0.39.jar
- tomcat-juli-9.0.39.jar
- tomcat-servlet-api-9.0.39.jar
- tomcat-util-9.0.39.jar
- tomcat-util-scan-9.0.39.jar
- univocity-parsers-2.7.3.jar
- validation-api-1.1.0.Final.jar
- velocity-1.7.jar
- woodstox-core-5.0.3.jar
- xbean-asm6-shaded-4.8.jar
- xercesImpl-2.12.0.jar
- xml-apis-1.4.01.jar
- xz-1.0.jar
- zjsonpatch-0.3.0.jar

- zookeeper-3.5.6-ei-302002.jar
- zookeeper-jute-3.5.6-ei-302002.jar
- zstd-jni-1.4.9-1.jar

Spark 2.3.2 Dependencies

- accessors-smart-1.2.jar
- activation-1.1.1.jar
- aircompressor-0.8.jar
- alluxio-2.3.1-luxor-SNAPSHOT-client.jar
- antlr-2.7.7.jar
- antlr4-runtime-4.8-1.jar
- antlr-runtime-3.4.jar
- aopalliance-1.0.jar
- aopalliance-repackaged-2.4.0-b34.jar
- apache-log4j-extras-1.2.17.jar
- arpack_combined_all-0.1.jar
- arrow-format-0.8.0.jar
- arrow-memory-0.8.0.jar
- arrow-vector-0.8.0.jar
- asm-5.0.4.jar
- audience-annotations-0.5.0.jar
- automaton-1.11-8.jar
- avro-1.7.7.jar
- avro-ipc-1.7.7.jar
- avro-ipc-1.7.7-tests.jar
- avro-mapred-1.7.7-hadoop2.jar
- java-sdk-bundle-1.11.271.jar
- base64-2.3.8.jar
- bcprov-jdk15on-1.66.jar
- bcprov-jdk15on-1.66.jar
- bonecp-0.8.0.RELEASE.jar
- breeze_2.11-0.13.2.jar
- breeze-macros_2.11-0.13.2.jar
- calcite-avatica-1.2.0-incubating.jar
- calcite-core-1.2.0-incubating.jar
- calcite-linq4j-1.2.0-incubating.jar
- checker-qual-2.11.1.jar
- chill_2.11-0.8.4.jar
- chill-java-0.8.4.jar
- commons-beanutils-1.9.4.jar

- commons-cli-1.2.jar
- commons-codec-2.0-20130428.202122-59.jar
- commons-collections-3.2.2.jar
- commons-collections4-4.2.jar
- commons-compiler-3.0.8.jar
- commons-compress-1.4.1.jar
- commons-configuration2-2.1.1.jar
- commons-crypto-1.0.0-20191105.jar
- commons-daemon-1.0.13.jar
- commons-dbcp-1.4.jar
- commons-dbcp2-2.7.0.jar
- commons-httpclient-3.1.jar
- commons-io-2.5.jar
- commons-lang-2.6.jar
- commons-lang3-3.5.jar
- commons-logging-1.2.jar
- commons-math3-3.4.1.jar
- commons-net-2.2.jar
- commons-pool-1.5.4.jar
- commons-pool2-2.8.0.jar
- commons-text-1.3.jar
- compress-lzf-1.0.3.jar
- core-1.1.2.jar
- curator-client-4.2.0.jar
- curator-framework-4.2.0.jar
- curator-recipes-2.7.1.jar
- datanucleus-api-jdo-3.2.6.jar
- datanucleus-core-3.2.10.jar
- datanucleus-rdbms-3.2.9.jar
- derby-10.12.1.1.jar
- dnsjava-2.1.7.jar
- ehcache-3.3.1.jar
- eigenbase-properties-1.1.5.jar
- error_prone_annotations-2.3.4.jar
- failureaccess-1.0.1.jar
- fastutil-8.2.3.jar
- ffmpeg-4.3.1-1.5.4.jar
- ffmpeg-4.3.1-1.5.4-linux-x86_64.jar
- flatbuffers-1.2.0-3f79e055.jar
- generex-1.0.2.jar

- geronimo-jcache_1.0_spec-1.0-alpha-1.jar
- gson-2.2.4.jar
- guava-29.0-jre.jar
- guice-4.0.jar
- guice-servlet-4.0.jar
- hadoop-annotations-3.1.1-ei-302002.jar
- hadoop-auth-3.1.1-ei-302002.jar
- hadoop-3.1.1-ei-302002.jar
- hadoop-client-3.1.1-ei-302002.jar
- hadoop-common-3.1.1-ei-302002.jar
- hadoop-hdfs-3.1.1-ei-302002.jar
- hadoop-hdfs-client-3.1.1-ei-302002.jar
- hadoop-3.1.1-41.jar
- hadoop-mapreduce-client-common-3.1.1-ei-302002.jar
- hadoop-mapreduce-client-core-3.1.1-ei-302002.jar
- hadoop-mapreduce-client-jobclient-3.1.1-ei-302002.jar
- hadoop-minikdc-3.1.1-ei-302002.jar
- hadoop-yarn-api-3.1.1-ei-302002.jar
- hadoop-yarn-client-3.1.1-ei-302002.jar
- hadoop-yarn-common-3.1.1-ei-302002.jar
- hadoop-yarn-registry-3.1.1-ei-302002.jar
- hadoop-yarn-server-common-3.1.1-ei-302002.jar
- hadoop-yarn-server-web-proxy-3.1.1-ei-302002.jar
- hamcrest-core-1.3.jar
- HikariCP-java7-2.4.12.jar
- hive-common-1.2.1-2.1.0.dli-20201111.064115-91.jar
- hive-exec-1.2.1-2.1.0.dli-20201111.064444-91.jar
- hive-metastore-1.2.1-2.1.0.dli-20201111.064230-91.jar
- hk2-api-2.4.0-b34.jar
- hk2-locator-2.4.0-b34.jar
- hk2-utils-2.4.0-b34.jar
- hppc-0.7.2.jar
- htrace-core4-4.2.0-incubating-1.0.0.jar
- httpclient-4.5.4.jar
- httpcore-4.4.7.jar
- ivy-2.4.0.jar
- j2objc-annotations-1.3.jar
- jackson-annotations-2.10.0.jar
- jackson-core-2.10.0.jar
- jackson-core-asl-1.9.13-atlassian-4.jar

- jackson-databind-2.10.0.jar
- jackson-dataformat-yaml-2.10.0.jar
- jackson-datatype-jsr310-2.10.3.jar
- jackson-jaxrs-base-2.10.3.jar
- jackson-jaxrs-json-provider-2.10.3.jar
- jackson-mapper-asl-1.9.13-atlassian-4.jar
- jackson-module-jaxb-annotations-2.10.3.jar
- jackson-module-paranamer-2.10.0.jar
- jackson-module-scala_2.11-2.10.0.jar
- jakarta.activation-api-1.2.1.jar
- jakarta.xml.bind-api-2.3.2.jar
- janino-3.0.8.jar
- javacpp-1.5.4.jar
- javacpp-1.5.4-linux-x86_64.jar
- javacv-1.5.4.jar
- JavaEWAH-1.1.7.jar
- javassist-3.18.1-GA.jar
- javax.annotation-api-1.2.jar
- javax.inject-1.jar
- javax.inject-2.4.0-b34.jar
- javax.servlet-api-3.1.0.jar
- javax.ws.rs-api-2.0.1.jar
- java-xmlbuilder-1.1.jar
- javolution-5.3.1.jar
- jaxb-api-2.2.11.jar
- jcip-annotations-1.0-1.jar
- jcl-over-slf4j-1.7.26.jar
- jdo-api-3.0.1.jar
- jersey-client-2.23.1.jar
- jersey-common-2.23.1.jar
- jersey-container-servlet-2.23.1.jar
- jersey-container-servlet-core-2.23.1.jar
- jersey-guava-2.23.1.jar
- jersey-media-jaxb-2.23.1.jar
- jersey-server-2.23.1.jar
- jets3t-0.9.4.jar
- jetty-http-9.4.31.v20200723.jar
- jetty-io-9.4.31.v20200723.jar
- jetty-security-9.4.31.v20200723.jar
- jetty-server-9.4.31.v20200723.jar

- jetty-servlet-9.4.31.v20200723.jar
- jetty-util-9.4.31.v20200723.jar
- jetty-util-ajax-9.4.31.v20200723.jar
- jetty-webapp-9.4.31.v20200723.jar
- jetty-xml-9.4.31.v20200723.jar
- joda-time-2.9.3.jar
- jodd-core-4.2.0.jar
- json-20200518.jar
- json4s-ast_2.11-3.2.11.jar
- json4s-core_2.11-3.2.11.jar
- json4s-jackson_2.11-3.2.11.jar
- json-sanitizer-1.2.1.jar
- json-smart-2.3.jar
- jsp-api-2.1.jar
- jsr305-3.0.2.jar
- jta-1.1.jar
- jtransforms-2.4.0.jar
- jul-to-slf4j-1.7.26.jar
- junit-4.11.jar
- kerb-admin-1.0.1.jar
- kerb-client-1.0.1.jar
- kerb-common-1.0.1.jar
- kerb-core-1.0.1.jar
- kerb-crypto-1.0.1.jar
- kerb-identity-1.0.1.jar
- kerb-server-1.0.1.jar
- kerb-simplekdc-1.0.1.jar
- kerb-util-1.0.1.jar
- kerby-asn1-1.0.1.jar
- kerby-config-1.0.1.jar
- kerby-pkix-1.0.1.jar
- kerby-util-1.0.1.jar
- kerby-xdr-1.0.1.jar
- kryo-shaded-3.0.3.jar
- kubernetes-client-4.9.2-20200804.jar
- kubernetes-model-4.9.2-20200804.jar
- kubernetes-model-common-4.9.2-20200804.jar
- leveledbjni-all-1.8-20191105.jar
- libfb303-0.9.3.jar
- libthrift-0.12.0.jar

- listenablefuture-9999.0-empty-to-avoid-conflict-with-guava.jar
- log4j-1.2.17-cloudera1.jar
- log4j-rolling-appender-20131024-2017.jar
- logging-interceptor-3.14.4.jar
- luxor-encrypt-2.1.0-20201106.065437-53.jar
- luxor-fs3-2.1.0-20201106.065612-53.jar
- luxor-obs-fs3-2.1.0-20201106.065616-53.jar
- luxor-rpc_2.11-2.1.0-20201106.065541-53.jar
- luxor-rpc-protobuf2-2.1.0-20201106.065551-53.jar
- lz4-java-1.7.1.jar
- machinist_2.11-0.6.1.jar
- macro-compat_2.11-1.1.1.jar
- metrics-core-3.1.5.jar
- metrics-graphite-3.1.5.jar
- metrics-jmx-4.1.12.1.jar
- metrics-json-3.1.5.jar
- metrics-jvm-3.1.5.jar
- minlog-1.3.0.jar
- mssql-jdbc-6.2.1.jre7.jar
- netty-3.10.6.Final.jar
- netty-all-4.1.51.Final.jar
- nimbus-jose-jwt-8.19.jar
- objenesis-2.1.jar
- okhttp-3.14.4.jar
- okio-1.17.2.jar
- opencsv-2.3.jar
- opencsv-4.6.jar
- opencv-4.3.0-2.jar
- orc-core-1.4.4-nohive.jar
- orc-mapreduce-1.4.4-nohive.jar
- oro-2.0.8.jar
- osgi-resource-locator-1.0.1.jar
- paranamer-2.8.jar
- parquet-column-1.8.3.jar
- parquet-common-1.8.3.jar
- parquet-encoding-1.8.3.jar
- parquet-format-2.3.1.jar
- parquet-hadoop-1.8.3.jar
- parquet-hadoop-bundle-1.6.0.jar
- parquet-jackson-1.8.3.jar

- parquet-format-2.3.1.jar
- parquet-hadoop-1.8.3.jar
- parquet-hadoop-bundle-1.6.0.jar
- parquet-jackson-1.8.3.jar
- postgresql-42.2.14.jar
- protobuf-java-2.5.0.jar
- py4j-0.10.7.jar
- pyrolite-4.13.jar
- re2j-1.1.jar
- RoaringBitmap-0.5.11.jar
- scala-compiler-2.11.12.jar
- scala-library-2.11.12.jar
- scalap-2.11.0.jar
- scala-parser-combinators_2.11-1.1.0.jar
- scala-reflect-2.11.12.jar
- scala-xml_2.11-1.0.5.jar
- secComponentApi-1.0.5c.jar
- shapeless_2.11-2.3.2.jar
- slf4j-api-1.7.30.jar
- slf4j-log4j12-1.7.30.jar
- snakeyaml-1.24.jar
- snappy-java-1.1.7.5.jar
- spark-catalyst_2.11-2.3.2.0101-2.1.0.dli-20201111.073826-143.jar
- spark-core_2.11-2.3.2.0101-.0.dli-20201111.073836-134.jar
- spark-graphx_2.11-2.3.2.0101-2.1.0.dli-20201111.073847-129.jar
- spark-hive_2.11-2.3.2.0101-.0.dli-20201111.073854-132.jar
- spark-kubernetes_2.11-2.3.2.0101-2.1.0.dli-20201111.073916-85.jar
- spark-kvstore_2.11-2.3.2.0101-2.1.0.dli-20201111.073933-127.jar
- spark-launcher_2.11-2.3.2.0101-2.1.0.dli-20201111.073940-127.jar
- spark-mllib_2.11-2.3.2.0101-2.1.0.dli-20201111.073946-127.jar
- spark-mllib-local_2.11-2.3.2.0101-2.1.0.dli-20201111.073953-127.jar
- spark-network-common_2.11-2.3.2.0101-2.1.0.dli-20201111.073959-127.jar
- spark-network-shuffle_2.11-2.3.2.0101-2.1.0.dli-20201111.074007-127.jar
- spark-om_2.11-2.3.2.0101-.0.dli-20201111.074019-125.jar
- spark-repl_2.11-2.3.2.0101-2.1.0.dli-20201111.074028-125.jar
- spark-sketch_2.11-2.3.2.0101-2.1.0.dli-20201111.074035-125.jar
- spark-sql_2.11-2.3.2.0101-2.1.0.dli-20201111.074041-126.jar
- spark-streaming_2.11-2.3.2.0101-2.1.0.dli-20201111.074100-123.jar
- spark-tags_2.11-2.3.2.0101-2.1.0.dli-20201111.074136-123.jar
- spark-tags_2.11-2.3.2.0101-2.1.0.dli-20201111.074141-124-tests.jar

- spark-unsafe_2.11-2.3.2.0101-2.1.0.dli-20201111.074144-123.jar
- spark-uquery_2.11-2.3.2.0101-2.1.0.dli-20201111.074906-210.jar
- spark-yarn_2.11-2.3.2.0101-2.1.0.dli-20201111.074151-123.jar
- spire_2.11-0.13.0.jar
- spire-macros_2.11-0.13.0.jar
- ST4-4.3.1.jar
- stax2-api-3.1.4.jar
- stax-api-1.0-2.jar
- stream-2.7.0.jar
- stringtemplate-3.2.1.jar
- token-provider-1.0.1.jar
- univocity-parsers-2.5.9.jar
- validation-api-1.1.0.Final.jar
- woodstox-core-5.0.3.jar
- xbean-asm5-shaded-4.4.jar
- xercesImpl-2.12.0.jar
- xml-apis-1.4.01.jar
- xz-1.0.jar
- zjsonpatch-0.3.0.jar
- zookeeper-3.5.6-ei-302002.jar
- zookeeper-jute-3.5.6-ei-302002.jar
- zstd-jni-1.4.4-11.jar

Flink 1.7.2 Dependencies

For details about the sample code of a Flink 1.7.2 program, see **luxor-demo\dli-flink-demo** in [DLI examples](#).

- bcpkix-jdk15on-1.60.jar
- bcprov-jdk15on-1.60.jar
- commons-codec-1.9.jar
- commons-configuration-1.7.jar
- deeplearning4j-core-0.9.1.jar
- deeplearning4j-nlp-0.9.1.jar
- deeplearning4j-nn-0.9.1.jar
- ejml-cdense-0.33.jar
- ejml-core-0.33.jar
- ejml-ddense-0.33.jar
- ejml-dsparse-0.33.jar
- ejml-experimental-0.33.jar
- ejml-fdense-0.33.jar
- ejml-simple-0.33.jar

- ejml-zdense-0.33.jar
- elsa-3.0.0-M7.jar
- esdk-obs-java-3.1.3.jar
- flink-cep_2.11-1.7.0.jar
- flink-cep-scala_2.11-1.7.0.jar
- flink-dist_2.11-1.7.0.jar
- flink-gelly_2.11-1.7.0.jar
- flink-gelly-scala_2.11-1.7.0.jar
- flink-ml_2.11-1.7.0.jar
- flink-python_2.11-1.7.0.jar
- flink-queryable-state-runtime_2.11-1.7.0.jar
- flink-shaded-curator-1.7.0.jar
- flink-shaded-hadoop2-uber-1.7.0.jar
- flink-table_2.11-1.7.0.jar
- guava-26.0-jre.jar
- hadoop-3.1.1-41-20201014.085840-4.jar
- httpasyncclient-4.1.2.jar
- httpclient-4.5.12.jar
- httpcore-4.4.4.jar
- httpcore-nio-4.4.4.jar
- java-xmlbuilder-1.1.jar
- jna-4.1.0.jar
- libtensorflow-1.12.0.jar
- log4j-api-2.16.0.jar
- log4j-core-2.16.0.jar
- log4j-over-slf4j-1.7.21.jar
- logback-classic-1.2.3.jar
- logback-core-1.2.3.jar
- nd4j-api-0.9.1.jar
- nd4j-native-0.9.1.jar
- nd4j-native-api-0.9.1.jar
- nd4j-native-platform-0.9.1.jar
- okhttp-3.14.8.jar
- okio-1.14.0.jar
- slf4j-api-1.7.21.jar
- tensorflow-1.12.0.jar

Flink 1.10 Dependencies

For details about the sample code for a Flink 1.10 program, see [Writing Data to OBS Using Flink Jar](#).

Only queues created after December 2020 can use the Flink 1.10 dependencies.

- bcpkix-jdk15on-1.60.jar
- bcprov-jdk15on-1.60.jar
- commons-codec-1.9.jar
- commons-configuration-1.7.jar
- deeplearning4j-core-0.9.1.jar
- deeplearning4j-nlp-0.9.1.jar
- deeplearning4j-nn-0.9.1.jar
- ejml-cdense-0.33.jar
- ejml-core-0.33.jar
- ejml-ddense-0.33.jar
- ejml-dsparse-0.33.jar
- ejml-experimental-0.33.jar
- ejml-fdense-0.33.jar
- ejml-simple-0.33.jar
- ejml-zdense-0.33.jar
- elsa-3.0.0-M7.jar
- esdk-obs-java-3.20.6.1.jar
- flink-cep_2.11-1.10.0.jar
- flink-cep-scala_2.11-1.10.0.jar
- flink-dist_2.11-1.10.0.jar
- flink-python_2.11-1.10.0.jar
- flink-queryable-state-runtime_2.11-1.10.0.jar
- flink-sql-client_2.11-1.10.0.jar
- flink-state-processor-api_2.11-1.10.0.jar
- flink-table_2.11-1.10.0.jar
- flink-table-blink_2.11-1.10.0.jar
- guava-26.0-jre.jar
- hadoop-3.1.1-41.jar
- httpasyncclient-4.1.2.jar
- httpclient-4.5.3.jar
- httpcore-4.4.4.jar
- httpcore-nio-4.4.4.jar
- java-xmlbuilder-1.1.jar
- jna-4.1.0.jar
- libtensorflow-1.12.0.jar
- log4j-over-slf4j-1.7.26.jar
- logback-classic-1.2.3.jar
- logback-core-1.2.3.jar
- nd4j-api-0.9.1.jar
- nd4j-native-0.9.1.jar

- nd4j-native-api-0.9.1.jar
- nd4j-native-platform-0.9.1.jar
- okhttp-3.14.8.jar
- okio-1.14.0.jar
- secComponentApi-1.0.5.jar
- slf4j-api-1.7.26.jar
- tensorflow-1.12.0.jar

Flink 1.12 Dependencies

- bcpkix-jdk15on-1.60.jar
- bcprov-jdk15on-1.60.jar
- clickhouse-jdbc-0.3.1-ei-313001-SNAPSHOT.jar
- commons-codec-1.9.jar
- commons-configuration-1.7.jar
- dataflow-fs-obs-2.0.0-20220226.034402-190.jar
- deeplearning4j-core-0.9.1.jar
- deeplearning4j-nlp-0.9.1.jar
- deeplearning4j-nn-0.9.1.jar
- ejml-cdense-0.33.jar
- ejml-core-0.33.jar
- ejml-ddense-0.33.jar
- ejml-dsparse-0.33.jar
- ejml-experimental-0.33.jar
- ejml-fdense-0.33.jar
- ejml-simple-0.33.jar
- ejml-zdense-0.33.jar
- elsa-3.0.0-M7.jar
- flink-changelog-json-1.12.2-ei-313001-dli-2022011002.jar
- flink-csv-1.12.2-ei-313001-dli-2022011002.jar
- flink-dist_2.11-1.12.2-ei-313001-dli-2022011002.jar
- flink-json-1.12.2-ei-313001-dli-2022011002.jar
- flink-kubernetes_2.11-1.12.2-ei-313001-dli-2022011002.jar
- flink-metrics-prometheus_2.11-1.12.2-ei-313001-dli-2022011002.jar
- flink-obs-hadoop-fs-2.0.0-20220226.034421-73.jar
- flink-s3-fs-hadoop-1.12.2.jar
- flink-shaded-zookeeper-3.6.3-ei-313001-SNAPSHOT.jar
- flink-sql-avro-1.12.2-ei-313001-dli-2022011002.jar
- flink-sql-avro-confluent-registry-1.12.2-ei-313001-dli-2022011002.jar
- flink-table_2.11-1.12.2-ei-313001-dli-2022011002.jar
- flink-table-blink_2.11-1.12.2-ei-313001-dli-2022011002.jar

- guava-18.0.jar
- guava-26.0-jre.jar
- hadoop-hdfs-client-3.1.1-ei-302002.jar
- hadoop-3.1.1-46.jar
- hadoop-plugins-8.1.3-313001-SNAPSHOT.jar
- httpasyncclient-4.1.2.jar
- httpclient-4.5.3.jar
- httpcore-4.4.4.jar
- httpcore-nio-4.4.4.jar
- java-xmlbuilder-1.1.jar
- jna-4.1.0.jar
- libtensorflow-1.12.0.jar
- log4j-1.2-api-2.17.1.jar
- log4j-api-2.17.1.jar
- log4j-core-2.17.1.jar
- log4j-slf4j-impl-2.17.1.jar
- luxor-encrypt-2.0.0-20220405.072004-199.jar
- luxor-fs3-2.0.0-20220405.072025-195.jar
- luxor-obs-fs3-2.0.0-20220405.072030-195.jar
- manager-hadoop-security-crypter-8.1.3-313001-SNAPSHOT.jar
- manager-wc2frm-8.1.3-313001-SNAPSHOT.jar
- mrs-obs-provider-3.1.1.49.jar
- nd4j-api-0.9.1.jar
- nd4j-native-0.9.1.jar
- nd4j-native-api-0.9.1.jar
- nd4j-native-platform-0.9.1.jar
- okhttp-3.14.8.jar
- okio-1.14.0.jar
- ranger-obs-client-0.1.1.jar
- secComponentApi-1.0.5.jar
- slf4j-api-1.7.26.jar
- tensorflow-1.12.0.jar

6 Job Templates

6.1 Managing SQL Templates

To facilitate SQL operation execution, DLI allows you to customize query templates or save the SQL statements in use as templates. After templates are saved, you do not need to compile SQL statements. Instead, you can directly perform the SQL operations using the templates.

SQL templates include sample templates and custom templates. The default sample template contains 22 standard TPC-H query statements, which can meet most TPC-H test requirements. For details, see [TPC-H Sample Data in the SQL Template](#).

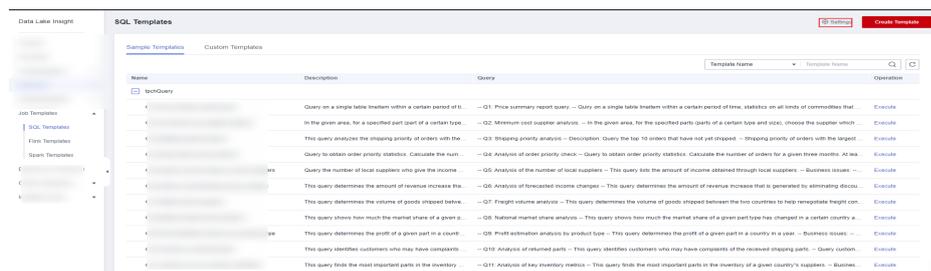
SQL template management provides the following functions:

- [Sample Templates](#)
- [Custom Templates](#)
- [Creating a Template](#)
- [Executing the Template](#)
- [Searching for a Template](#)
- [Modifying a Template](#)
- [Deleting a Template](#)

Table Settings

In the upper right corner of the **SQL Template** page, click **Set Property** to determine whether to display templates by group.

Figure 6-1 Template settings



If you select **Display by Group**, the following display modes are available:

- Expand the first group
- Expand all
- Collapse All

Sample Templates

The current sample template contains 22 standard TPC-H query statements. You can view the template name, description, and statements. For details about TPC-H examples, see [TPC-H Sample Data in the SQL Template](#).

Table 6-1 Template management parameters

Parameter	Description
Name	Indicates the template name. <ul style="list-style-type: none"> ● A template name can contain only digits, letters, and underscores (_), but cannot start with an underscore (_) or contain only digits. It cannot be left empty. ● The template name can contain a maximum of 50 characters.
Description	Description of the template you create.
Statement	SQL statement created as the template.
Operation	Execute: After you click this button, the system switches to the SQL Editor page, where you can modify or directly perform the statement as required. For details, see Executing the Template .

The existing sample templates apply to the following scenarios:

- Price summary report query
- Minimum cost supplier analysis

- Shipping priority analysis
- Analysis of order priority check
- Analysis of the number of local suppliers
- Analysis of forecasted income changes
- Freight volume analysis
- National market share analysis
- Profit estimation analysis by product type
- Analysis of returned parts
- Analysis of key inventory indicators
- Freight mode and command priority analysis
- Consumer allocation analysis
- Promotion effect analysis
- Analysis of the supplier with the largest contribution
- Analysis of the relationship between parts and suppliers
- Revenue analysis of small-lot orders
- Customer analysis for large orders
- Discounted revenue analysis
- Analysis of potential component improvements
- Analysis of suppliers who fail to deliver goods on time
- Global sales opportunity analysis

Custom Templates

The custom template list displays all templates you have created. You can view the template name, description, statements, and more.

Table 6-2 Template management parameters

Parameter	Description
Name	Indicates the template name. <ul style="list-style-type: none">• A template name can contain only digits, letters, and underscores (_), but cannot start with an underscore (_) or contain only digits. It cannot be left empty.• The template name can contain a maximum of 50 characters.
Description	Description of the template you create.
Statement	SQL statement created as the template.

Parameter	Description
Operation	<ul style="list-style-type: none"> • Execute: After you click this button, the system switches to the SQL Editor page, where you can modify or directly perform the statement as required. For details, see Executing the Template. • Modify: Click Modify. In the displayed Modify Template dialog box, modify the template information as required. For details, see Modifying a Template.

Creating a Template

You can create a template on either the **Job Templates** or the **SQL Editor** page.

- To create a template on the **Job Templates** page:
 - a. On the left of the management console, choose **Job Templates > SQL Templates**.
 - b. On the **SQL Templates** page, click **Create Template** to create a template.

Enter the template name, SQL statement, and description information. For details, see [Table 6-3](#).

Figure 6-2 Creating a template

Create Template

You can create 89 more templates. [Increase quota](#).

* Name

* Statement
0/10,000

Description
0/256

Group Use Existing Use new Do not use

* Group Name

OK Cancel

Table 6-3 Parameter description

Parameter	Description
Name	Indicates the template name. <ul style="list-style-type: none">• A template name can contain only digits, letters, and underscores (_), but cannot start with an underscore (_) or contain only digits. It cannot be left empty.• The template name can contain a maximum of 50 characters.
Statement	SQL statement to be saved as a template.
Description	Description of the template you create.
Group	<ul style="list-style-type: none">• Use existing• Use new• Do not use
Group Name	If you select Use existing or Use new , you need to enter the group name.

- c. Click **OK**.
- To create a template on the **SQL Editor** page:
 - a. On the left of the management console, click **SQL Editor**.
 - b. In the SQL job editing area of the displayed **SQL Editor** page, click **More** in the upper right corner and choose **Save as Template**.
Enter the template name, SQL statement, and description information.
For details, see [Table 6-3](#).
 - c. Click **OK**.

Executing the Template

Perform the template operation as follows:

1. On the left of the management console, choose **Job Templates > SQL Templates**.
2. On the **SQL Templates** page, select a template and click **Execute** in the **Operation** column. The **SQL Editor** page is displayed, and the corresponding SQL statement is automatically entered in the SQL job editing window.
3. In the upper right corner of the SQL job editing window, Click **Execute** to run the SQL statement. After the execution is complete, you can view the execution result below the current SQL job editing window.

Searching for a Template

On the **SQL Templates** page, you can enter the template name keyword in the search box on the upper right corner to search for the desired template.

Modifying a Template

Only custom templates can be modified. To modify a template, perform the following steps:

- Step 1** On the **SQL Templates** page, locate the target template and click **Modify** in the **Operation** column.
- Step 2** In the displayed **Modify Template** dialog box, modify the template name, statement, and description as required.
- Step 3** Click **OK**.

----End

Deleting a Template

On the **SQL Templates** page, select one or more templates to be deleted and click **Delete** to delete the selected templates.

6.2 Managing Flink Templates

Flink templates include sample templates and custom templates. You can modify an existing sample template to meet the actual job logic requirements and save time for editing SQL statements. You can also customize a job template based on your habits and methods so that you can directly invoke or modify the template in later jobs.

Flink template management provides the following functions:

- [Flink SQL Sample Template](#)
- [Custom Templates](#)
- [Creating a Template](#)
- [Creating a Job Based on a Template](#)
- [Modifying a Template](#)
- [Deleting a Template](#)

Flink SQL Sample Template

The template list displays existing sample templates for Flink SQL jobs. [Table 1](#) describes the parameters in the template list.

The scenarios of sample templates can be different, which are subject to the console.

Table 6-4 Parameters in the Flink SQL sample template list

Parameter	Description
Name	Name of a template, which has 1 to 64 characters and only contains letters, digits, hyphens (-), and underlines (_).
Description	Description of a template. It contains 0 to 512 characters.

Parameter	Description
Operation	Create Job: Create a job directly by using the template. After a job is created, the system switches to the Edit page under Job Management .

Custom Templates

The custom template list displays all Jar job templates. [Table 1](#) describes parameters in the custom template list.

Table 6-5 Parameters in the custom template list

Parameter	Description
Name	Name of a template, which has 1 to 64 characters and only contains letters, digits, hyphens (-), and underlines (_).
Description	Description of a template. It contains 0 to 512 characters.
Created	Time when a template is created.
Updated	Latest time when a template is modified.
Operation	<ul style="list-style-type: none">• Edit: Modify a template that has been created.• Create Job: Create a job directly by using the template. After a job is created, the system switches to the Edit page under Job Management.• More:<ul style="list-style-type: none">- Delete: Delete a created template.- Tags: View or add tags.

Creating a Template

You can create a template using any of the following methods:

- Creating a template on the **Template Management** page
 - a. In the left navigation pane of the DLI management console, choose **Job Templates > Flink Templates**.
 - b. Click **Create Template** in the upper right corner of the page. The **Create Template** dialog box is displayed.
 - c. Specify **Name** and **Description**.

Figure 6-3 Creating a Flink template

Create Template
×

Type Flink SQL ▼

★ Name Enter a name.

Description Description

Tags

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#) ↻

To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

Table 6-6 Template parameters

Parameter	Description
Name	Name of a template, which has 1 to 64 characters and only contains letters, digits, hyphens (-), and underlines (_). NOTE The template name must be unique.
Description	Description of a template. It contains 0 to 512 characters.

Parameter	Description
Tags	<p>Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS).</p> <p>For details, see Tag Management Service User Guide.</p> <p>NOTE</p> <ul style="list-style-type: none"> • A maximum of 20 tags can be added. • Only one tag value can be added to a tag key. • The key name in each resource must be unique. • Tag key: Enter a tag key name in the text box. <p>NOTE</p> <p>A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed, but the value cannot start or end with a space or start with <code>_sys_</code>.</p> <ul style="list-style-type: none"> • Tag value: Enter a tag value in the text box. <p>NOTE</p> <p>A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:=-@) are allowed. The value cannot start or end with a space.</p>

d. Click **OK** to enter the editing page.

The [Table 6-7](#) describes the parameters on the template editing page.

Table 6-7 Template parameters

Parameter	Description
Name	You can modify the template name.
Description	You can modify the template description.
Saving Mode	<ul style="list-style-type: none"> • Save Here: Save the modification to the current template. • Save as New: Save the modification as a new template.
SQL statement editing area	In the area, you can enter detailed SQL statements to implement business logic. For details about how to compile SQL statements, see Data Lake Insight SQL Syntax Reference .
Save	Save the modifications.
Create Job	Use the current template to create a job.
Format	Format SQL statements. After SQL statements are formatted, you need to compile SQL statements again.

Parameter	Description
Theme Settings	Change the font size, word wrap, and page style (black or white background).

- e. In the SQL statement editing area, enter SQL statements to implement service logic. For details about how to compile SQL statements, see [Data Lake Insight SQL Syntax Reference](#).
 - f. After the SQL statement is edited, click **Save** in the upper right corner to complete the template creation.
 - g. (Optional) If you do not need to modify the template, click **Create Job** in the upper right corner to create a job based on the current template. For details about how to create a job, see [Creating a Flink SQL Job](#), and [Creating a Flink Jar Job](#).
- Creating a template based on an existing job template
 - a. In the left navigation pane of the DLI management console, choose **Job Templates > Flink Templates**. Click the **Custom Templates** tab.
 - b. In the row where the desired template is located in the custom template list, click **Edit** under **Operation** to enter the **Edit** page.
 - c. After the modification is complete, set **Saving Mode** to **Save as New**.
 - d. Click **Save** in the upper right corner to save the template as a new one.
 - Creating a template using a created job
 - a. In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
 - b. Click **Create Job** in the upper right corner. The **Create Job** page is displayed.
 - c. Specify parameters as required.
 - d. Click **OK** to enter the editing page.
 - e. After the SQL statement is compiled, click **Set as Template**.
 - f. In the **Set as Template** dialog box that is displayed, specify **Name** and **Description** and click **OK**.
 - Creating a template based on the existing job
 - a. In the left navigation pane of the DLI management console, choose **Job Management > Flink Jobs**. The **Flink Jobs** page is displayed.
 - b. In the job list, locate the row where the job that you want to set as a template resides, and click **Edit** in the **Operation** column.
 - c. After the SQL statement is compiled, click **Set as Template**.
 - d. In the **Set as Template** dialog box that is displayed, specify **Name** and **Description** and click **OK**.

Creating a Job Based on a Template

You can create jobs based on sample templates or custom templates.

1. In the left navigation pane of the DLI management console, choose **Job Templates > Flink Templates**.

2. In the sample template list, click **Create Job** in the **Operation** column of the target template. For details about how to create a job, see [Creating a Flink SQL Job](#) and [Creating a Flink Jar Job](#).

Modifying a Template

After creating a custom template, you can modify it as required. The sample template cannot be modified, but you can view the template details.

1. In the left navigation pane of the DLI management console, choose **Job Templates > Flink Templates**. Click the **Custom Templates** tab.
2. In the row where the template you want to modify is located in the custom template list, click **Edit** in the **Operation** column to enter the **Edit** page.
3. In the SQL statement editing area, modify the SQL statements as required.
4. Set **Saving Mode** to **Save Here**.
5. Click **Save** in the upper right corner to save the modification.

Deleting a Template

You can delete a custom template as required. The sample templates cannot be deleted. Deleted templates cannot be restored. Exercise caution when performing this operation.

1. In the left navigation pane of the DLI management console, choose **Job Templates > Flink Templates**. Click the **Custom Templates** tab.
2. In the custom template list, select the templates you want to delete and click **Delete** in the upper left of the custom template list.
Alternatively, you can delete a template by performing the following operations: In the custom template list, locate the row where the template you want to delete resides, and click **More > Delete** in the **Operation** column.
3. In the displayed dialog box, click **Yes**.

6.3 Managing Spark SQL Templates

Scenario

You can modify a sample template to meet the Spark job requirements, saving time for editing SQL statements.

Currently, the cloud platform does not provide preset Spark templates. You can customize Spark job templates. This section describes how to create a Spark job template.

Creating a Management Job Template

To create a Spark job template, you can save a Spark job information as a template.

1. In the left navigation pane of the DLI management console, choose **Job Templates > Spark Templates**. The **Spark Templates** page is displayed.

2. Configure job parameters by referring to [Creating a Spark Job](#).
3. After you finish editing the job, click **Save as Template**.
4. Enter the template name and other information as you need.
5. Set a template group for future management.
6. Click **OK**.

6.4 Appendix

6.4.1 TPC-H Sample Data in the SQL Template

TPC-H Sample Data

TPC-H is a test set developed by the Transaction Processing Performance Council (TPC) to simulate decision-making support applications. It is widely used in academia and industry to evaluate the performance of decision-making support technology. This business test has higher requirements on vendors, because it can comprehensively evaluate the overall business computing capability. With universal business significance, is widely used in analysis of bank credit, credit card, telecom operation, tax, as well as tobacco industry decision-making analysis.

The TPC-H benchmark test is developed from TPC-D (a standard specified by TPC in 1994 and used as the test benchmark for decision-making support systems). TPC-H implements a 3NF data warehouse that contains eight basic relationships, with a data volume range from 1 GB to 3 TB. The TPC-H benchmark test includes 22 queries (Q1 to Q22). The main evaluation indicator is the response time of each query (from submission to result return). The unit of the TPC-H benchmark test is the query number per hour (QphH@size). **H** indicates the average number of complex queries per hour. **size** indicates the size of database, which reflects the query processing capability of the system. TPC-H can evaluate key performance parameters that other tests cannot evaluate, because it is modeled based on the actual production and operation environment. In a word, the TPC-H standard by TPC meets the test requirements of data warehouse and motivate vendors and research institutes to stretch the limit of this technology.

In this example, DLI directly queries the TPC-H dataset on OBS. DLI has generated a standard TPC-H-2.18 dataset of 100 MB which is uploaded to the tpch folder on OBS. The read-only permission is granted to you to facilitate query operations.

TPC-H Test and Metrics

TPC-H test is divided into three sub-tests: data loading test, Power test, and Throughput test. Data loading indicates the process of setting up a test database, and the loading test is to test the data loading ability of DBMS. The first test is data loading test that tests data loading time, which is time-consuming. The second test is Power test, also called raw query. After data loading test is complete, the database is in the initial state without any other operation, especially the data in the buffer is not tested. Power test requires that the 22 queries be executed once in sequence and a pair of RF1 and RF2 operations be executed at the same time. The third test is Throughput test, the core and most complex test, more similar to the actual application environment. With multiple

query statement groups and a pair of RF1 and RF2 update flows, Throughput test pose greater pressure on the SUT system than Power test does.

The basic data in the test is related to the execution time (the time of each data loading step, each query execution, and each update execution), based on which you can calculate the data loading time, Power@Size, Throughput@Size, qphH@Size and \$/QphH@Size.

Power@Size is the result of the Power test, which is defined as the reciprocal of the geometric average value of the query time and change time. The formula is as follows:

$$\text{TPC-H Power@Size} = \frac{3600 * SF}{\sqrt[24]{\prod_{i=1}^{i=22} QI(i,0) * \prod_{j=1}^{j=2} RI(j,0)}}$$

Size indicates the data size. SF is the scaling factor of data scale. QI (i, 0) indicates the time of the ith query, in seconds. R (I j, 0) is the update time of RFj, in seconds.

Throughput@Size is the Throughput test result, which is defined as the reciprocal of the average value of all query execution time. The formula is as follows:

$$\text{QphH@Size} = \sqrt{\text{Power @ Size} * \text{Throughput @ Size}}$$

Service Scenario

You can use the built-in TPC-H test suite of DLI to perform interactive query without uploading data.

Advantages of DLI Built-in TPC-H

- You can log in to DLI and get permission to run SQL statements without creating tables or import data.
- The 22 preset TPC-H SQL query templates with rich functions meet the requirements of most business scenarios. You do not need to download TPC-H query statements, which saves your time and energy.
- Data Lake gives you brand-new experience of serverless DLI product within the minimum time.

Precautions

When a sub-account uses the TPC-H test suite, the main account needs to grant the sub-account the OBS access permission and the permission to view the main account table. If the master account has not logged in to DLI, the sub-account needs to have the permissions to create databases and tables in addition to the preceding permissions.

Operation Description

For details, see [Managing SQL Templates](#).

7 Enhanced Datasource Connections

7.1 Overview

What Is Enhanced Datasource Connection?

Typically, you cannot use DLI to directly access a data source in a VPC other than the one where DLI is because the network between DLI and the data source is disconnected. For proper access, you need to establish a network connection between them.

DLI provides enhanced connections. Establishing a VPC peering connection allows DLI to communicate with the VPC of the data source, supporting cross-source data analysis.

For details about the data sources that support cross-source access, see [Cross-Source Analysis Development Methods](#).

CAUTION

In cross-source development scenarios, there is a risk of password leakage if datasource authentication information is directly configured. You are advised to use the datasource authentication provided by DLI. For details, see [Introduction](#).

Constraints

- Datasource connections cannot be created for the **default** queue.
- Flink jobs can directly access DIS, OBS, and SMN data sources without using datasource connections.
- Enhanced connections can only be created for yearly/monthly and pay-per-use queues.
- **VPC Administrator** permissions are required for enhanced connections to use VPCs, subnets, routes, VPC peering connections.

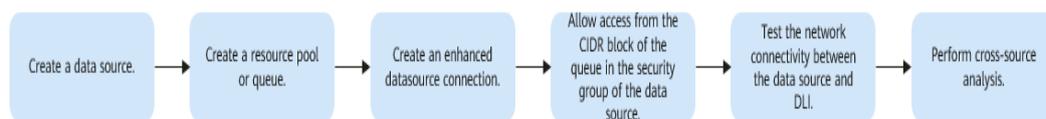
You can set these permissions by referring to [Service Authorization](#).

- If you use an enhanced datasource connection, the CIDR block of the elastic resource pool or queue cannot overlap with that of the data source.
- Only queues bound with datasource connections can access datasource tables.
- Datasource tables do not support the preview function.
- When checking the connectivity of datasource connections, the constraints on IP addresses are as follows:
 - The IP address must be valid, which consists of four decimal numbers separated by periods (.). The value ranges from 0 to 255.
 - During the test, you can add a port after the IP address and separate them with colons (:). The port can contain a maximum of five digits. The value ranges from 0 to 65535.
For example, **192.168.xx.xx** or **192.168.xx.xx:8181**.
- When checking the connectivity of datasource connections, the constraints on domain names are as follows:
 - The domain name can contain 1 to 255 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.
 - The top-level domain name must contain at least two letters, for example, **.com**, **.net**, and **.cn**.
 - During the test, you can add a port after the domain name and separate them with colons (:). The port can contain a maximum of five digits. The value ranges from 0 to 65535.
For example, **example.com:8080**.

Cross-Source Analysis Process

To use DLI for cross-source analysis, you need to create a datasource connection to connect DLI to the data source, and then develop jobs to access the data source.

Figure 7-1 Cross-source analysis flowchart



Helpful Links

Creation Method	Enhanced Datasource Connection
Console	Creating an Enhanced Datasource Connection
API	Creating an Enhanced Datasource Connection

7.2 Cross-Source Analysis Development Methods

Cross-Source Analysis

If DLI needs to access external data sources, you need to establish enhanced datasource connections to enable the network between DLI and the data sources, and then develop different types of jobs to access the data sources. This is the process of DLI cross-source analysis.

This section describes how to develop data sources supported by DLI for cross-source analysis.

Notes

- Flink jobs can directly access DIS, OBS, and SMN data sources without using datasource connections.
- You are advised to use enhanced datasource connections to connect DLI to data sources.

Cross-Source Analysis Development Methods

[Table 7-1](#) lists the data sources supported by DLI and the corresponding development methods.

Table 7-1 Syntax reference for cross-source analysis

Service	Spark SQL Job	Spark Jar Job	Flink OpenSource SQL Job	Flink Jar Job
CloudTable HBase	<ul style="list-style-type: none"> • Create an HBase association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	<ul style="list-style-type: none"> • HBase source table • HBase result table • HBase dimension table 	-
CloudTable OpenTSDB	<ul style="list-style-type: none"> • Create an OpenTSDB association table • Insert data • Insert data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	-	-

Service	Spark SQL Job	Spark Jar Job	Flink OpenSource SQL Job	Flink Jar Job
CSS	<ul style="list-style-type: none"> • Create a CSS association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	<ul style="list-style-type: none"> • Elasticsearch result table 	-
DCS Redis	<ul style="list-style-type: none"> • Create a DCS association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code • Java sample code 	<ul style="list-style-type: none"> • Redis source table • Redis result table • Redis dimension table 	Flink job sample
DDS Mongo	<ul style="list-style-type: none"> • Create a DDS association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code • Java sample code 	-	-
DMS	-	-	<ul style="list-style-type: none"> • Kafka source table • Kafka result table 	-
GaussDB(DWS)	<ul style="list-style-type: none"> • Create a GaussDB(DWS) association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	<ul style="list-style-type: none"> • GaussDB(DWS) source table • GaussDB(DWS) result table • GaussDB(DWS) dimension table 	Flink job sample

Service	Spark SQL Job	Spark Jar Job	Flink OpenSource SQL Job	Flink Jar Job
MRS HBase	<ul style="list-style-type: none"> • Create an HBase association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	<ul style="list-style-type: none"> • HBase source table • HBase result table • HBase dimension table 	Flink job sample
MRS Kafka	-	-	<ul style="list-style-type: none"> • Kafka source table • Kafka result table 	Flink job sample
MRS OpenTSDB	<ul style="list-style-type: none"> • Create an OpenTSDB association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	-	-
RDS MySQL	<ul style="list-style-type: none"> • Create an RDS association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	<ul style="list-style-type: none"> • MySQL CDC source table 	-
RDS PostGre	<ul style="list-style-type: none"> • Create an RDS association table • Insert data • Query data 	<ul style="list-style-type: none"> • Scala sample code • PySpark sample code 	<ul style="list-style-type: none"> • Postgres CDC source table 	-

7.3 Creating an Enhanced Datasource Connection

Scenario

Create an enhanced datasource connection for DLI to access, import, query, and analyze data of other data sources.

For example, to connect DLI to the MRS, RDS, CSS, Kafka, or GaussDB(DWS) data source, you need to enable the network between DLI and the VPC of the data source.

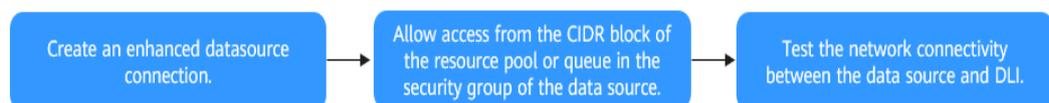
Create an enhanced datasource connection on the console.

Constraints

- Datasource connections cannot be created for the **default** queue.
- Flink jobs can directly access DIS, OBS, and SMN data sources without using datasource connections.
- Enhanced connections can only be created for yearly/monthly and pay-per-use queues.
- **VPC Administrator** permissions are required for enhanced connections to use VPCs, subnets, routes, VPC peering connections.
You can set these permissions by referring to [Service Authorization](#).
- If you use an enhanced datasource connection, the CIDR block of the elastic resource pool or queue cannot overlap with that of the data source.
- Only queues bound with datasource connections can access datasource tables.
- Datasource tables do not support the preview function.
- When checking the connectivity of datasource connections, the constraints on IP addresses are as follows:
 - The IP address must be valid, which consists of four decimal numbers separated by periods (.). The value ranges from 0 to 255.
 - During the test, you can add a port after the IP address and separate them with colons (:). The port can contain a maximum of five digits. The value ranges from 0 to 65535.
For example, **192.168.xx.xx** or **192.168.xx.xx:8181**.
- When checking the connectivity of datasource connections, the constraints on domain names are as follows:
 - The domain name can contain 1 to 255 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.
 - The top-level domain name must contain at least two letters, for example, **.com**, **.net**, and **.cn**.
 - During the test, you can add a port after the domain name and separate them with colons (:). The port can contain a maximum of five digits. The value ranges from 0 to 65535.
For example, **example.com:8080**.

Process

Figure 7-2 Enhanced datasource connection creation flowchart



Prerequisites

- An elastic resource pool or queue has been created.
- You have obtained the VPC, subnet, private IP address, port, and security group information of the external data source.
- The security group of the external data source has allowed access from the CIDR block of the elastic resource pool or queue.

Procedure

Step 1 Create an Enhanced Datasource Connection

1. Log in to the DLI management console.
2. In the left navigation pane, choose **Datasource Connections**.
3. On the **Enhanced** tab page displayed, click **Create**.

Configure parameters according to [Table 7-2](#).

Table 7-2 Parameters

Parameter	Description
Connection Name	Name of the created datasource connection. <ul style="list-style-type: none">- The name can contain only letters, digits, and underscores (_). The parameter must be specified.- A maximum of 64 characters are allowed.
Resource Pool	It binds an elastic resource pool or queue that uses a datasource connection. This parameter is optional. Only dedicated queues charged in yearly/monthly or pay-per-use billing mode can be bound to elastic resource pools. In regions where this function is available, an elastic resource pool with the same name is created by default for the queue created in "Creating a Queue." NOTE Before using an enhanced datasource connection, you must bind a queue and ensure that the VPC peering connection is in the Active state.
VPC	VPC used by the data source.
Subnet	Subnet used by the data source.
Route Table	Route table of the subnet. NOTE <ul style="list-style-type: none">- The route table is associated with the subnet used by the destination data source, which is not the table containing the route you add by Manage Route in the Operation column. The route you add on the Manage Route page is contained in the route table associated with the subnet used by the queue to be bound.- The subnet used by the destination data source must be different from that used by the queue to be bound. Otherwise, a segment conflict occurs.

Parameter	Description
Host Information	<p>In this text field, you can configure the mapping between host IP addresses and domain names so that jobs can only use the configured domain names to access corresponding hosts. This parameter is optional.</p> <p>For example, when accessing the HBase cluster of MRS, you need to configure the host name (domain name) and IP address of the ZooKeeper instance. Enter one record in each line in the format of <i>IP address Host name/Domain name</i>.</p> <p>Example:</p> <pre>192.168.0.22 node-masterxxx1.com 192.168.0.23 node-masterxxx2.com</pre> <p>For details about how to obtain host information, see How Do I Obtain MRS Host Information?.</p>
Tags	<p>Tags used to identify cloud resources. A tag includes the tag key and tag value. If you want to use the same tag to identify multiple cloud resources, that is, to select the same tag from the drop-down list box for all services, you are advised to create predefined tags on the Tag Management Service (TMS).</p> <p>For details, see Tag Management Service User Guide.</p> <p>NOTE</p> <ul style="list-style-type: none"> - A maximum of 20 tags can be added. - Only one tag value can be added to a tag key. - The key name in each resource must be unique. <p>- Tag key: Enter a tag key name in the text box.</p> <p>NOTE</p> <p>A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (<code>_.:=-@</code>) are allowed, but the value cannot start or end with a space or start with <code>_sys_</code>.</p> <p>- Tag value: Enter a tag value in the text box.</p> <p>NOTE</p> <p>A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (<code>_.:=-@</code>) are allowed. The value cannot start or end with a space.</p>

4. Click **OK**.

After the creation is complete, the enhanced datasource connection is in the **Active** state, indicating that the connection is successfully created.

Step 2 Security Group Where the Data Source Belongs Allows Access from the CIDR Block of the Elastic Resource Pool

1. On the DLI management console, obtain the network segment of the elastic resource pool or queue.

- Choose **Resources** > **Queue Management** from the left navigation pane. On the page displayed, locate the queue on which jobs are running, and click the button next to the queue name to obtain the CIDR block of the queue.
2. Log in to the VPC console and find the VPC the data source belongs to.
 3. On the network console, choose **Virtual Private Cloud** > **Network Interfaces**. On the **Network Interfaces** tab page displayed, search for the security group name, click **More** in the **Operation** column, and select **Change Security Group**.
 4. In the navigation pane on the left, choose **Access Control** > **Security Groups**.
 5. Click the name of the security group to which the external data source belongs.
 6. Click the **Inbound Rules** tab and add a rule to allow access from the CIDR block of the queue. See [Figure 7-3](#).

Configure the inbound rule parameters according to [Table 7-3](#).

Figure 7-3 Adding an inbound rule

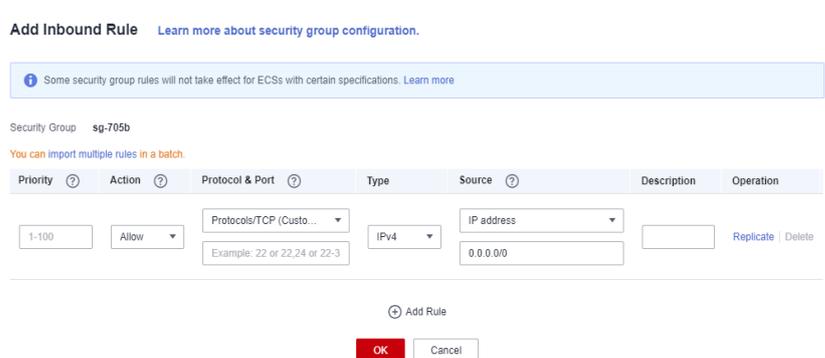


Table 7-3 Inbound rule parameters

Parameter	Description	Example Value
Priority	Priority of a security group rule. The priority value ranges from 1 to 100. The default value is 1, indicating the highest priority. A smaller value indicates a higher priority of a security group rule.	1
Action	Action of the security group rule.	Allow

Parameter	Description	Example Value
Protocol & Port	<ul style="list-style-type: none">– Network protocol. The value can be All, TCP, UDP, ICMP, or GRE.– Port: Port or port range over which the traffic can reach your instance. The port ranges from 1 to 65535.	In this example, select TCP . Leave the port blank or set it to the data source port.
Type	Type of IP addresses.	IPV4
Source	Allows access from IP addresses or instances in another security group.	In this example, enter the obtained queue CIDR block.
Description	Supplementary information about the security group rule. This parameter is optional.	–

Step 3 Test the Connectivity Between the DLI Queue and the Data Source

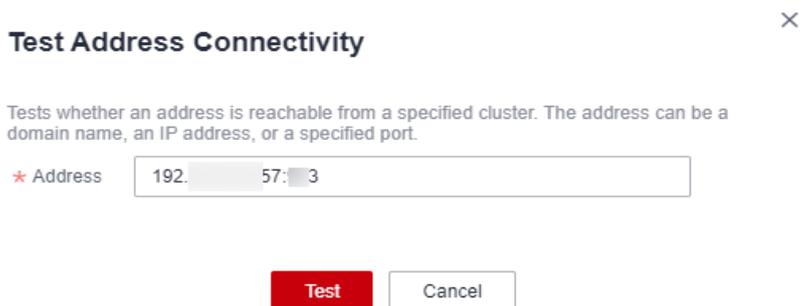
1. Obtain the private IP address and port number of the data source.
Take the RDS data source as an example. On the **Instances** page, click the target DB instance. On the page displayed, locate the **Connection Information** pane and view the private IP address. In the **Connection Information** pane, locate the **Database Port** to view the port number of the RDS DB instance.
2. In the navigation pane of the DLI management console, choose **Resources > Queue Management**.
3. Locate the queue bound with the enhanced datasource connection, click **More** in the **Operation** column, and select **Test Address Connectivity**.
4. Enter the data source connection address and port number to test the network connectivity.

Format: *IP address:Port number*

CAUTION

Before testing the connection, ensure that the security group of the external data source has allowed access from the CIDR block of the queue.

Figure 7-4 Testing the network connectivity between the queue and the data source



----End

Related Operations

- [Why Is a Datasource Connection Successfully Created But the Network Connectivity Test Fails?](#)
- [Configuring the Connection Between a DLI Queue and a Data Source in a Private Network](#)
- [Configuring the Connection Between a DLI Queue and a Data Source in the Internet](#)

7.4 Deleting an Enhanced Datasource Connection

Scenario

Delete an enhanced datasource connection that is no longer used on the console.

Procedure

1. Log in to the DLI management console.
2. In the left navigation pane, choose **Datasource Connections**.
3. On the **Enhanced** tab page displayed, locate the enhanced datasource connection to be deleted and click **Delete** in the **Operation** column.
4. In the dialog box displayed, click **Yes**.

7.5 Modifying Host Information

Scenario

Host information is the mapping between host IP addresses and domain names. After you configure host information, jobs can only use the configured domain names to access corresponding hosts. After a datasource connection is created, you can modify the host information.

When accessing the HBase cluster of MRS, you need to configure the host name (domain name) and IP address of the instance.

Constraints

You have obtained the MRS host information by referring to [How Do I Obtain MRS Host Information?](#)

Modifying Host Information

1. Log in to the DLI management console.
2. In the left navigation pane, choose **Datasource Connections**.
3. On the **Enhanced** tab page displayed, locate the enhanced datasource connection to be modified, click **More** in the **Operation** column, and select **Modify Host**.
4. In the **Modify Host** dialog box displayed, enter the obtained host information. Enter host information in the format of *Host IP address Host name*. Information about multiple hosts is separated by line breaks.

Example:

```
192.168.0.22 node-masterxxx1.com
```

```
192.168.0.23 node-masterxxx2.com
```

Obtain the MRS host information by referring to [How Do I Obtain MRS Host Information?](#)

5. Click **OK**.

How Do I Obtain MRS Host Information?

- **Method 1: View MRS host information on the management console.**

To obtain the host name and IP address of an MRS cluster, for example, MRS 3.x, perform the following operations:

- a. Log in to the MRS management console.
- b. On the **Active Clusters** page displayed, click your desired cluster to access its details page.
- c. Click the **Components** tab.
- d. Click **ZooKeeper**.
- e. Click the **Instance** tab to view corresponding service IP addresses. You can select any service IP address.
- f. Modify host information by referring to [Modifying Host Information](#).

NOTE

If the MRS cluster has multiple IP addresses, enter any service IP address when creating a datasource connection.

- **Method 2: Obtain MRS host information from the /etc/hosts file on an MRS node.**

- a. Log in to any MRS node as user **root**.
- b. Run the following command to obtain MRS hosts information. Copy and save the information.

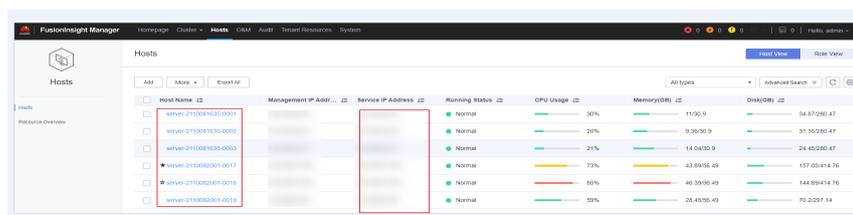
cat /etc/hosts

Figure 7-5 Obtaining hosts information

```
[root@node-master1nqt1 ~]# cat /etc/hosts
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
10.10.10.10 hadoop hadoop.com
10.10.10.10 manager
192.168.0.44 node-master1nqt1.b27fd346-a6fb-42ef-bcc0-d964639baafb.com node-master1nqt1.b27fd346-a6fb-42ef-bcc0-d964639baafb.com.
192.168.0.198 node-master2Y8cb.b27fd346-a6fb-42ef-bcc0-d964639baafb.com node-master2Y8cb.b27fd346-a6fb-42ef-bcc0-d964639baafb.com.
192.168.0.188 node-ana-corezyNL.b27fd346-a6fb-42ef-bcc0-d964639baafb.com node-ana-corezyNL.b27fd346-a6fb-42ef-bcc0-d964639baafb.com.
192.168.0.136 node-ana-corezyKB.b27fd346-a6fb-42ef-bcc0-d964639baafb.com node-ana-corezyKB.b27fd346-a6fb-42ef-bcc0-d964639baafb.com.
192.168.0.174 node-ana-coreVPY0.b27fd346-a6fb-42ef-bcc0-d964639baafb.com node-ana-coreVPY0.b27fd346-a6fb-42ef-bcc0-d964639baafb.com.
[root@node-master1nqt1 ~]#
[root@node-master1nqt1 ~]#
[root@node-master1nqt1 ~]#
```

- c. Modify host information by referring to [Modifying Host Information](#).
- **Method 3: Log in to FusionInsight Manager to obtain host information.**
 - a. Log in to FusionInsight Manager.
 - b. On FusionInsight Manager, click **Hosts**. On the **Hosts** page, obtain the host names and service IP addresses of the MRS hosts.

Figure 7-6 FusionInsight Manager



- c. Modify host information by referring to [Modifying Host Information](#).

7.6 Binding and Unbinding a Queue

Constraints

- The CIDR block of the DLI queue that is bound with a datasource connection cannot overlap with that of the data source.
- The default queue cannot be bound with a connection.

Binding a Queue

Before using an enhanced datasource connection, you must bind a queue and ensure that the VPC peering connection is in the **Active** state.

- **Method 1**
In the **Enhanced** tab, select a connection and click **Bind Queue** in the **Operation** column. In the dialog box that is displayed, select the queues to be bound and click **OK**.
- **Method 2**
Click the name of the selected connection. The **Details** page is displayed. Click **Create** in the upper left corner. In the dialog box that is displayed, select the queues to be bound and click **OK**.

Viewing Details about a Bound Queue

On the **Enhanced** tab page, select a connection and click the connection name to view the information about the bound queue.

Table 7-4 Parameters in the details list of datasource connection queues

Parameter	Description
VPC Peering ID	ID of the VPC peering connection created in the cluster to which the queue belongs. NOTE A VPC peering connection is created for each queue bound to an enhanced datasource connection. The VPC peering connection is used for cross-VPC communication. Ensure that the security group used by the data source allows access from the DLI queue CIDR block, and do not delete the VPC peering connection during the datasource connection.
Name	Name of a bound queue.
Connection Status	Datasource connection status. The following three statuses are available: <ul style="list-style-type: none">• Creating• Active• Failed NOTE If the connection status is Failed , click  on the left to view the detailed error information.
Updated	Time when a connection is updated. The connections in the connection list can be displayed according to the update time in ascending or descending order.
Operation	Unbind Queue: This operation is used to unbind a datasource connection from a queue.

Unbinding a Queue

If you do not need to use an enhanced datasource connection, you can unbind the queue from it to release resources.

- Method 1

On the **Enhanced** tab page, select a connection and click **More > Unbind Queue** in the **Operation** column. In the dialog box that is displayed, select the queues to be unbound and click **OK**.

- Method 2

Click the name of the target connection. The **Details** page is displayed. Select the queue to be unbound and click **Unbind Queue** in the **Operation** column. In the displayed dialog box, confirm the queue to be unbound and click **Yes**.

7.7 Adding a Route

Scenario

A route is configured with the destination, next hop type, and next hop to determine where the network traffic is directed. Routes are classified into system routes and custom routes.

After an enhanced connection is created, the subnet is automatically associated with the default route. You can add custom routes as needed to forward traffic destined for the destination to the specified next hop.

NOTE

- When an enhanced connection is created, the associated route table is the one associated with the subnet of the data source.
- The route to be added in the **Add Route** dialog box must be one in the route table associated with the subnet of the resource pool.
- The subnet of the data source must be different from that used by the resource pool. Otherwise, a network segment conflict occurs.

Procedure

1. Log in to the DLI management console.
2. In the left navigation pane, choose **Datasource Connections**.
3. On the **Enhanced** tab page displayed, locate the row containing the enhanced connection to which a route needs to be added, and add the route.
 - Method 1:
 - i. On the **Enhanced** tab page displayed, locate the enhanced datasource connection to which a route needs to be added and click **Manage Route** in the **Operation** column.
 - ii. Click **Add Route**.
 - iii. In the **Add Route** dialog box, enter the route information. For details about the parameters, see [Table 7-5](#).
 - iv. Click **OK**.
 - Method 2:
 - i. On the **Enhanced** tab page displayed, locate the enhanced datasource connection to which a route needs to be added, click **More** in the **Operation** column, and select **Add Route**.
 - ii. In the **Add Route** dialog box, enter the route information. For details about the parameters, see [Table 7-5](#).
 - iii. Click **OK**.

Table 7-5 Parameters for adding a custom route

Parameter	Description
Route Name	Name of a custom route, which is unique in the same enhanced datasource scenario. The name can contain 1 to 64 characters. Only digits, letters, underscores (_), and hyphens (-) are allowed.
IP Address	Custom route CIDR block. The CIDR block of different routes can overlap but cannot be the same.

4. After adding a route, you can view the route information on the route details page.

7.8 Deleting a Route

Scenario

Delete a route that is no longer used.

Constraints

A custom route table cannot be deleted if it is associated with a subnet.

Procedure

1. Log in to the DLI management console.
2. In the left navigation pane, choose **Datasource Connections**.
3. On the **Enhanced** tab page displayed, locate the row containing the enhanced connection from which the route needs to be deleted, and delete the route.
 - Method 1:
 - i. On the **Enhanced** tab page displayed, locate the enhanced connection from which the route needs to be deleted and click **Manage Route** in the **Operation** column.
 - ii. Locate the route to be deleted and click **Delete** in the **Operation** column.
 - iii. In the dialog box displayed, click **OK**.
 - Method 2:
 - i. On the **Enhanced** tab page displayed, locate the enhanced connection from which the route needs to be deleted, click **More** in the **Operation** column, and select **Delete Route**.
 - ii. In the **Delete Route** dialog box displayed, confirm the route information.
 - iii. Click **Yes**.

7.9 Enhanced Connection Permission Management

Scenario

Enhanced connections support user authorization by project. After authorization, users in the project have the permission to perform operations on the enhanced connection, including viewing the enhanced connection, binding a created resource pool to the enhanced connection, and creating custom routes. In this way, the enhanced connection can be used across projects. Grant and revoke permissions to and from a user for an enhanced connection.

NOTE

- If the authorized projects belong to different users in the same region, you can use the user account of the authorized projects to log in.
- If the authorized projects belong to the same user in the same region, you can use the current account to switch to the corresponding project.

Use Cases

Project B needs to access the data source of project A. The operations are as follows:

- For Project A:
 - a. Log in to DLI using the account of project A.
 - b. Create an enhanced datasource connection **ds** in DLI based on the VPC information of the corresponding data source.
 - c. Grant project B the permission to access the enhanced datasource connection **ds**.
- For Project B:
 - a. Log in to DLI using the account of project B.
 - b. Bind the enhanced datasource connection **ds** to a queue.
 - c. (Optional) Set host information and create a route.

After creating a VPC peering connection and route between the enhanced datasource connection of project A and the queue of project B, you can create a job in the queue of project B to access the data source of project A.

Procedure

1. Log in to the DLI management console.
2. In the left navigation pane, choose **Datasource Connections**.
3. On the **Enhanced** tab page displayed, locate the desired enhanced connection, click **More** in the **Operation** column, and select **Manage Permission**.
 - **Granting permission**
 - i. In the **Permissions** dialog box displayed, select **Grant Permission for Set Permission**.

- ii. Enter the project ID.
 - iii. Click **OK** to grant the resource pool operation permission to the project.
- **Revoking permission**
- i. In the **Permissions** dialog box displayed, select **Revoke Permission** for **Set Permission**.
 - ii. Select a project ID.
 - iii. Click **OK** to revoke the resource pool operation permission from the specified project.

7.10 Enhanced Datasource Connection Tag Management

Scenario

A tag is a key-value pair customized by users and used to identify cloud resources. It helps users to classify and search for cloud resources. A tag consists of a tag key and a tag value.

If you use tags in other cloud services, you are advised to create the same tag key-value pairs for cloud resources used by the same business to keep consistency.

DLI supports the following two types of tags:

- Resource tags: non-global tags created on DLI.
- Predefined tags: global tags created on Tag Management Service (TMS).
For more information about predefined tags, see [Tag Management Service User Guide](#).

DLI allows you to add, modify, or delete tags for datasource connections.

Procedure

1. In the left navigation pane of the DLI management console, choose **Datasource Connections**.
2. In the **Operation** column of the link, choose **More > Tags**.
3. The tag management page is displayed, showing the tag information about the current connection.
4. Click **Add/Edit Tag**. The **Add/Edit Tag** dialog is displayed. Add or edit tag keys and values and click **OK**.

Figure 7-7 Adding/Editing a tag

×

Add/Edit Tag

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#) ↻

To add a tag, enter a tag key and a tag value below.

Enter a tag key

Enter a tag value

Add

10 tags available for addition.

OK

Cancel

Table 7-6 Tag parameters

Parameter	Description
Tag key	<p>You can perform the following operations:</p> <ul style="list-style-type: none"> Click the text box and select a predefined tag key from the drop-down list. <p>To add a predefined tag, you need to create one on TMS and then select it from the Tag key drop-down list. You can click View predefined tags to go to the Predefined Tags page of the TMS console. Then, click Create Tag in the upper corner of the page to create a predefined tag.</p> <p>For details, see Creating Predefined Tags in <i>Tag Management Service User Guide</i>.</p> <ul style="list-style-type: none"> Enter a tag key in the text box. <p>NOTE A tag key can contain a maximum of 128 characters. Only letters, digits, spaces, and special characters (_.:+=-@) are allowed, but the value cannot start or end with a space or start with _sys_.</p>
Tag value	<p>You can perform the following operations:</p> <ul style="list-style-type: none"> Click the text box and select a predefined tag value from the drop-down list. Enter a tag value in the text box. <p>NOTE A tag value can contain a maximum of 225 characters. Only letters, digits, spaces, and special characters (_.:+=-@) are allowed. The value cannot start or end with a space.</p>

5. Click **OK**.
6. (Optional) To delete a tag, locate the row where the tag locates in the tag list and click **Delete** in the **Operation** column to delete the tag.

8 Datasource Authentication

8.1 Introduction

What Is Datasource Authentication?

When developing across multiple sources, it is not recommended to configure authentication information directly in a job as it can lead to password leakage. Instead, you are advised to use either Data Encryption Workshop (DEW) or datasource authentication provided by DLI to securely store data source authentication information.

- DEW is a comprehensive cloud-based encryption service that addresses data security, key security, and complex key management issues. You are advised to use DEW to store authentication information for data sources. For details, see [Data Encryption Workshop \(DEW\)](#).
- Datasource authentication is used to manage authentication information for accessing specified data sources. After datasource authentication is configured, you do not need to repeatedly configure data source authentication information in jobs, improving data source authentication security while enabling DLI to securely access data sources.

This section describes how to use datasource authentication provided by DLI.

Constraints

- Only Spark SQL and Flink OpenSource SQL 1.12 jobs support datasource authentication.
- Flink jobs can use datasource authentication only on queues created after May 1, 2023.
- DLI supports four types of datasource authentication. Select an authentication type specific to each data source.
 - CSS: applies to 6.5.4 or later CSS clusters with the security mode enabled.
 - Kerberos: applies to MRS security clusters with Kerberos authentication enabled.
 - Kafka_SSL: applies to Kafka with SSL enabled.

- Password: applies to GaussDB(DWS), RDS, DDS, and DCS.

Datasource Authentication Types

DLI supports four types of datasource authentication. Select an authentication type specific to each data source.

- CSS: applies to 6.5.4 or later CSS clusters with the security mode enabled. During the configuration, you need to specify the username, password, and authentication certificate of the cluster and store the information in DLI through datasource authentication so that DLI can securely access CSS data sources. For details, see [Creating a CSS Datasource Authentication](#).
- Kerberos: applies to MRS security clusters with Kerberos authentication enabled. During the configuration, you need to specify MRS cluster authentication credentials, including the `krb5.conf` and `user.keytab` files. For details, see [Creating a Kerberos Datasource Authentication](#).
- Kafka_SSL: applies to Kafka with SSL enabled. During the configuration, you need to specify the KafkaTruststore path and password. For details, see [Creating a Kafka_SSL Datasource Authentication](#).
- Password: applies to GaussDB(DWS), RDS, DDS, and DCS data sources. During the configuration, you need to store the passwords of the data sources in DLI. For details, see [Creating a Password Datasource Authentication](#).

Jobs That Can Connect to Data Sources Through Datasource Authentication

Different types of jobs can connect to data sources through different types of datasource authentication.

- For details about the data sources that Spark SQL jobs can connect to through datasource authentication and their constraints, see [Table 8-1](#).
- For details about the data sources that Flink OpenSource SQL 1.12 jobs can connect to through datasource authentication and their constraints, see [Table 8-2](#).

Table 8-1 Data sources that Spark SQL jobs can connect to through datasource authentication

Datasource Authentication Type	Data Source	Constraints
CSS	CSS	The CSS cluster version must be 6.5.4 or later. The security mode has been enabled for the CSS cluster.
Password	GaussDB(DWS), RDS, DDS, and Redis	-

Table 8-2 Data sources that Flink OpenSource SQL 1.12 jobs can connect to through datasource authentication

Table Type	Datasource Authentication Type	Data Source	Constraints
Source table	Kerberos	HBase	Kerberos authentication has been enabled for the MRS cluster.
		Kafka	Kerberos authentication has been enabled for MRS Kafka.
	Kafka_SSL	Kafka	SASL_SSL authentication has been enabled for DMS Kafka. SASL authentication has been enabled for MRS Kafka. SSL authentication has been enabled for MRS Kafka.
	Password	GaussDB(DWS), RDS, and Redis	-
Result table	Kerberos	HBase	Kerberos authentication has been enabled for the MRS cluster.
		Kafka	Kerberos authentication has been enabled for MRS Kafka.
	Kafka_SSL	Kafka	SASL_SSL authentication has been enabled for DMS Kafka. SASL authentication has been enabled for MRS Kafka. SSL authentication has been enabled for MRS Kafka.
	Password	GaussDB(DWS), RDS, CSS, and Redis	-
Dimension table	Kerberos	HBase	Kerberos authentication has been enabled for the MRS cluster.
	Password	GaussDB(DWS), RDS, and Redis	-

8.2 Creating a CSS Datasource Authentication

Scenario

Create a CSS datasource authentication on the DLI console to store the authentication information of the CSS security cluster to DLI. This will allow you to

access to the CSS security cluster without having to configure a username and password in SQL jobs.

Create a datasource authentication for a CSS security cluster on the DLI console.

Notes

A CSS security cluster has been created and has met the following conditions:

- The cluster version is 6.5.4 or later.
- The security mode has been enabled for the cluster.

Procedure

1. Download the authentication credential of the CSS security cluster.
 - a. Log in to the CSS management console and choose **Clusters > Elasticsearch**.
 - b. On the **Clusters** page displayed, click the cluster name.
 - c. On the **Cluster Information** page displayed, find the security mode and download the certificate of the CSS security cluster.
2. Upload the authentication credential to the OBS bucket.
3. Create a datasource authentication.
 - a. Log in to the DLI management console.
 - b. Choose **Datasource Connections**. On the page displayed, click **Datasource Authentication**.
 - c. Click **Create**.

Configure CSS authentication parameters according to [Table 8-3](#).

Table 8-3 Parameters

Parameter	Description
Authentication Certificate	Name of the datasource authentication information to be created. <ul style="list-style-type: none">• The name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_).• The length of the database name cannot exceed 128 characters.• It is recommended that the name contain the CSS security cluster name to distinguish security authentication information of different clusters.
Type	Select CSS .
Username	Username for logging in to the security cluster.
Password	The password of the security cluster

Parameter	Description
Certificate Path	Enter the OBS path to which the security certificate is uploaded, that is, the OBS bucket address in 2 .

Figure 8-1 Creating a datasource authentication-CSS

Create Authentication

×

Type

* Authentication Certificate

* Username ?

* Password ?

* Certificate Path 📁

4. Create a table to access the CSS cluster.

When creating a table, associate the table with the created datasource authentication to access the CSS cluster.

For example, when using Spark SQL to create a table for accessing the CSS cluster, configure **es.certificate.name** to set the datasource authentication name and then connect to the CSS security cluster.

Use Spark SQL to create a table for accessing the CSS cluster by referring to [Creating a DLI Table and Associating It with CSS](#).

8.3 Creating a Kerberos Datasource Authentication

Scenario

Create a Kerberos datasource authentication on the DLI console to store the authentication information of the data source to DLI. This will allow you to access to the data source without having to configure a username and password in SQL jobs.

 NOTE

- When Kerberos authentication is enabled for MRS Kafka but SSL authentication is disabled, create a Kerberos authentication. When creating a table, configure **krb_auth_name** to associate the datasource authentication.
- If Kerberos authentication and SSL authentication are both enabled for MRS Kafka, you need to create Kerberos and Kafka_SSL authentications. When creating a table, configure **krb_auth_name** and **ssl_auth_name** to associate the datasource authentications.
- Datasource authentication is not required when Kerberos authentication is disabled but SASL authentication is enabled for MRS Kafka (for example, when a username and a password are used for PlainLoginModule authentication).
- When Kerberos authentication is disabled but SSL authentication is enabled for MRS Kafka, you need to create a Kafka_SSL authentication. When creating a table, configure **ssl_auth_name** to associate the datasource authentication.
- When Kerberos authentication is disabled but SASL authentication and SSL authentication are enabled for MRS Kafka, you need to create a Kafka_SSL authentication. When creating a table, configure **ssl_auth_name** to associate the datasource authentication.

Data Sources Supported by Kerberos Datasource Authentication

[Table 8-4](#) lists the data sources supported by Kerberos datasource authentication.

Table 8-4 Data sources supported by Kerberos datasource authentication

Job Type	Table Type	Data Source	Constraints
Flink OpenSource SQL	Source table	HBase	Kerberos authentication has been enabled for the MRS cluster.
		Kafka	Kerberos authentication has been enabled for MRS Kafka.
	Result table	HBase	Kerberos authentication has been enabled for the MRS cluster.
		Kafka	Kerberos authentication has been enabled for MRS Kafka.
	Dimension table	HBase	Kerberos authentication has been enabled for the MRS cluster.

Procedure

1. Download the authentication credential of the data source.
 - a. Log in to MRS Manager.
 - b. Choose **System > Permission > User**.
 - c. Click **More**, select **Download Authentication Credential**, save the file, and decompress it to obtain the **keytab** and **krb5.conf** files.

2. Upload the authentication credential to the OBS bucket.
3. Create a datasource authentication.
 - a. Log in to the DLI management console.
 - b. Choose **Datasource Connections**. On the page displayed, click **Datasource Authentication**.
 - c. Click **Create**.
Configure Kerberos authentication parameters according to [Table 8-5](#).

Table 8-5 Parameters

Parameter	Description
Type	Select Kerberos .
Authentication Certificate	Name of the datasource authentication to be created. <ul style="list-style-type: none">• The name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_).• The name can contain a maximum of 128 characters.• It is recommended that the name contain the MRS security cluster name to distinguish security authentication information of different clusters.
Username	Username for logging in to the security cluster.
krb5_conf Path	OBS path to which the krb5.conf file is uploaded. NOTE The renew_lifetime configuration item under [libdefaults] must be removed from krb5.conf . Otherwise, the "Message stream modified (41)" error may occur.
keytab Path	OBS path to which the user.keytab file is uploaded.

Figure 8-2 Creating a datasource authentication – Kerberos

4. Create a table to access the MRS cluster.
When creating a data source, associate the data source with the created datasource authentication to access the data source.

Table 8-6 lists the fields used to associate with the datasource authentication during table creation.

Table 8-6 Fields that are used to associate with Kerberos datasource authentication during table creation

Job Type	Data Source	Parameter	Mandatory	Data Type	Description
Flink OpenSource SQL	HBa se	krb_a uth_n ame	No	String	This field is used to associate datasource authentications when source, result, and dimension tables are created.

Job Type	Data Source	Parameter	Mandatory	Data Type	Description
	Kafka	krb_auth_name	No	String	<p>This field is used to associate datasource authentications when source and result tables are created.</p> <p>Name of the created Kerberos datasource authentication.</p> <p>If SASL_PLAINTEXT and Kerberos authentication are both used, you need to configure the following parameters:</p> <ul style="list-style-type: none"> • 'properties.sasl.mechanism' = 'GSSAPI' • 'properties.security.protocol' = 'SASL_PLAINTEXT'

For details about how to create a table, see *Data Lake Insight Syntax Reference*.

- Flink OpenSource SQL Syntax Reference: [Creating an HBase Source Table](#)

8.4 Creating a Kafka_SSL Datasource Authentication

Scenario

Create a Kafka_SSL datasource authentication on the DLI console to store the Kafka authentication information to DLI. This will allow you to access to Kafka instances without having to configure a username and password in SQL jobs.

 NOTE

- When Kerberos authentication is enabled for MRS Kafka but SSL authentication is disabled, create a Kerberos authentication. When creating a table, configure **krb_auth_name** to associate the datasource authentication.
- If Kerberos authentication and SSL authentication are both enabled for MRS Kafka, you need to create Kerberos and Kafka_SSL authentications. When creating a table, configure **krb_auth_name** and **ssl_auth_name** to associate the datasource authentications.
- Datasource authentication is not required when Kerberos authentication is disabled but SASL authentication is enabled for MRS Kafka (for example, when a username and a password are used for PlainLoginModule authentication).
- When Kerberos authentication is disabled but SSL authentication is enabled for MRS Kafka, you need to create a Kafka_SSL authentication. When creating a table, configure **ssl_auth_name** to associate the datasource authentication.
- When Kerberos authentication is disabled but SASL authentication and SSL authentication are enabled for MRS Kafka, you need to create a Kafka_SSL authentication. When creating a table, configure **ssl_auth_name** to associate the datasource authentication.

Data Sources Supported by Kafka_SSL Datasource Authentication

Table 8-7 lists the data sources supported by Kafka_SSL datasource authentication.

Table 8-7 Data sources supported by Kafka_SSL datasource authentication

Job Type	Table Type	Data Source	Constraints
Flink OpenSource SQL	Source table and result table	Kafka	SASL_SSL authentication has been enabled for DMS Kafka. SASL authentication has been enabled for MRS Kafka. SSL authentication has been enabled for MRS Kafka.

Procedure

1. Download the authentication credential.
 - **DMS Kafka**
 - i. Log in to the DMS (for Kafka) console and click a Kafka instance to access its details page.
 - ii. In the connection information, find the SSL certificate and click **Download**.
Decompress the downloaded **kafka-certs** package to obtain the **client.jks** and **phy_ca.crt** files.
 - **MRS Kafka**
 - i. Log in to MRS Manager.

- ii. Choose **System > Permission > User**.
 - iii. Click **More**, select **Download Authentication Credential**, save the file, and decompress it to obtain the truststore file.
 2. Upload the authentication credential to the OBS bucket.
 3. Create a datasource authentication.
 - a. Log in to the DLI management console.
 - b. Choose **Datasource Connections**. On the page displayed, click **Datasource Authentication**.
 - c. Click **Create**.
Configure Kafka authentication parameters according to [Table 8-8](#).

Table 8-8 Parameters

Parameter	Description
Type	Select Kafka_SSL .
Authentication Certificate	Name of the datasource authentication to be created. <ul style="list-style-type: none">• The name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_).• The name can contain a maximum of 128 characters.
Truststore Path	OBS path to which the SSL truststore file is uploaded. <ul style="list-style-type: none">• For MRS Kafka, enter the OBS path of the Truststore.jks file.• For DMS Kafka, enter the OBS path of the client.jks file.
Truststore Password	Truststore password.
Keystore Path	OBS path to which the SSL keystore file (key and certificate) is uploaded.
Keystore Password	Keystore (key and certificate) password.
Key Password	Password of the private key in the keystore file.

4. Access Kafka with SASL_SSL authentication enabled.
When creating a data source, associate the data source with the created datasource authentication to access the data source.
[Table 8-9](#) lists the fields used to associate with the datasource authentication during table creation.

Table 8-9 Fields that are used to associate with Kafka_SSL datasource authentication during table creation

Parameter	Mandatory	Data Type	Description
ssl_auth_name	No	String	<p>This field is used to associate datasource authentications when source, result, and dimension tables are created.</p> <p>Name of the created Kafka_SSL datasource authentication. This configuration is used when SSL is configured for Kafka.</p> <ul style="list-style-type: none">• If only SSL is used, configure the following parameter: 'properties.security.protocol' = 'SSL';• If SASL_SSL is used, configure the following parameters:<ul style="list-style-type: none">- 'properties.security.protocol' = 'SASL_SSL',- 'properties.sasl.mechanism' ='GSSAPI or PLAIN'- 'properties.sasl.jaas.config' = 'org.apache.kafka.common.security.plain.PlainLoginModule required username=\"xxx\" password=\"xxx\";'

For details about how to create a table, see *Data Lake Insight Syntax Reference*.

- Flink OpenSource SQL Syntax Reference: [Creating a Kafka Source Table](#)

8.5 Creating a Password Datasource Authentication

Scenario

Create a password datasource authentication on the DLI console to store passwords of the GaussDB(DWS), RDS, DCS, and DDS data sources to DLI. This will allow you to access to the data sources without having to configure a username and password in SQL jobs.

Data Sources Supported by Password Datasource Authentication

[Table 8-10](#) lists the data sources supported by password datasource authentication.

Table 8-10 Data sources supported by password datasource authentication

Job Type	Table Type	Data Source
Spark SQL	-	GaussDB(DWS), RDS, DDS, and Redis
Flink OpenSource SQL	Source table	GaussDB(DWS), RDS, and Redis
	Result table	GaussDB(DWS), RDS, CSS, and Redis
	Dimension table	GaussDB(DWS), RDS, and Redis

Procedure

1. Create a datasource authentication.
 - a. Log in to the DLI management console.
 - b. Choose **Datasource Connections**. On the page displayed, click **Datasource Authentication**.
 - c. Click **Create**.
Configure authentication parameters according to [Table 8-11](#).

Table 8-11 Parameters

Parameter	Description
Type	Select Password .
Authentication Certificate	Name of the datasource authentication to be created. <ul style="list-style-type: none">• The name can contain only digits, letters, and underscores (_), but cannot contain only digits or start with an underscore (_).• The name can contain a maximum of 128 characters.
Username	Username for accessing the data source.
Password	Password for accessing the data source.

Figure 8-3 Creating a datasource authentication – Password

✕

Create Authentication

Type

✱ Authentication Certificate

Username

?

✱ Password

👁 ?

2. Access the data source.

When creating a data source, associate the data source with the created datasource authentication to access the data source.

Table 8-12 lists the fields used to associate with the datasource authentication during table creation.

Table 8-12 Fields that are used to associate with password datasource authentication during table creation

Job Type	Parameter	Mandatory	Data Type	Description
Spark SQL	passwd auth	No	String	Name of datasource authentication. It is applicable to GaussDB(DWS), RDS, DDS, and Redis data sources.
Flink OpenSource SQL	pwd_auth_name	No	String	This field is used to associate datasource authentications when source, result, and dimension tables are created. Set pwd_auth_name to the name of the password datasource authentication. If this parameter is set, you do not need to configure a username and a password of the data source in SQL jobs.

For details about how to create a table, see *Data Lake Insight Syntax Reference*.

- Flink OpenSource SQL Syntax Reference: [Creating a GaussDB\(DWS\) Source Table](#)

8.6 Datasource Authentication Permission Management

Scenario

Grant permissions on a datasource authentication to users so multiple user jobs can use the datasource authentication without affecting each other.

Notes

- The administrator and the datasource authentication owner have all permissions. You do not need to set permissions for them, and their datasource authentication permissions cannot be modified by other users.
- When setting datasource authentication permissions for a new user, ensure that the user group to which the user belongs has the **Tenant Guest** permission.

For details about the **Tenant Guest** permission and how to apply for the permission, see [System Permissions](#) and [Creating a User Group](#) in *Identity and Access Management User Guide*.

Granting Permissions on Datasource Connections

1. Log in to the DLI management console.
2. Choose **Datasource Connections**. On the page displayed, click **Datasource Authentication**.
3. Locate the row containing the datasource authentication to be authorized and click **Manage Permission** in the **Operation** column. The **User Permissions** page is displayed.
4. Click **Grant Permission** in the upper right corner of the page. On the **Grant Permission** dialog box displayed, grant permissions on this datasource authentication to other users.

Figure 8-4 Granting permissions on datasource connections

Grant Permission

* Username

Select the permissions to be granted to the user

<input type="checkbox"/> Select all	<input type="checkbox"/> Update	<input type="checkbox"/> Delete
<input type="checkbox"/> Access	<input type="checkbox"/> Revoke Permission	<input type="checkbox"/> View Other User's Permissions
<input type="checkbox"/> Grant Permission		

Table 8-13 Permission granting parameters

Parameter	Description
Username	Name of the IAM user to whom permissions on the datasource connection are to be granted. NOTE The username is the name of an existing IAM user.
Select the permissions to be granted to the user	<ul style="list-style-type: none">• Access: This permission allows you to access the datasource connection.• Update: This permission allows you to update the datasource connection.• Delete: This permission allows you to delete the datasource connection.• Grant Permission: This permission allows you to grant the datasource connection permission to other users.• Grant Permission: This permission allows you to revoke the datasource connection permission to other users. However, you cannot revoke the permissions of the datasource connection owner.• View Other User's Permissions: This permission allows you to view the datasource connection permissions of other users.

Modifying the Permissions of Current User

1. Log in to the DLI management console.
2. Choose **Datasource Connections**. On the page displayed, click **Datasource Authentication**.
3. Locate the row containing the datasource authentication to be authorized and click **Manage Permission** in the **Operation** column. The **User Permissions** page is displayed.
4. Click **Set Permission** in the **Operation** column to modify the permissions of the current user. [Table 8-13](#) lists the detailed permission descriptions.

NOTE

- If all options under **Set Permission** are gray, you are not allowed to change permissions on this datasource connection. You can apply to the administrator, group owner, or other users who have the permission to grant permissions for the permissions to grant and revoke the datasource authentication permissions.
- To revoke all permissions of the current user, click **Revoke Permission** in the **Operation** column. The IAM user will no longer have any permission on the datasource authentication.

9 Global Configuration

9.1 Global Variables

What Is a Global Variable?

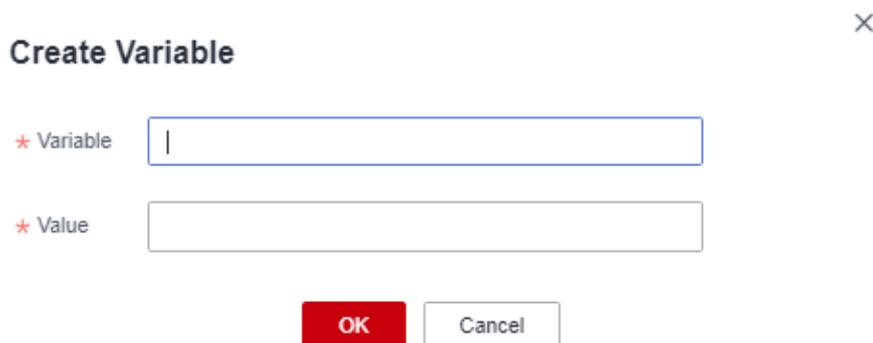
DLI allows you to set variables that are frequently used during job development as global variables on the DLI management console. This avoids repeated definitions during job editing and reduces development and maintenance costs. Global variables can be used to replace long and difficult variables, simplifying complex parameters and improving the readability of SQL statements.

This section describes how to create a global variable.

Creating Variables

1. In the navigation pane of the DLI console, choose **Global Configuration > Global Variables**.
2. On the **Global Variables** page, click **Create** in the upper right corner to create a global variable.

Figure 9-1 Creating a global variable



The screenshot shows a 'Create Variable' dialog box. It has a title bar with the text 'Create Variable' and a close button (X) in the top right corner. Below the title bar, there are two input fields. The first field is labeled '* Variable' and contains a vertical cursor. The second field is labeled '* Value' and is empty. At the bottom of the dialog, there are two buttons: 'OK' (a red button) and 'Cancel' (a white button with a grey border).

Table 9-1 Parameters description

Parameter	Description
Variable	Name of the created global variable.
Value	Global variable value.

- After creating a global variable, use **{{xxxx}}** in the SQL statement to replace the parameter value set as the global variable. **xxxx** indicates the variable name. For example, if you set global variable **abc** to represent the table name, replace the actual table name with **{{abc}}** in the table creation statement.

```
create table {{table_name}} (String1 String, int4 int, varchar1 varchar(10))  
partitioned by (int1 int,int2 int,int3 int)
```

NOTE

Do not use global variables in **OPTIONS** of the table creation statements.

Modifying Variables

On the **Global Variables** page, click **Modify** in the **Operation** column of a variable to modify the variable value.

NOTE

If there are multiple global variables with the same name in the same project under an account, delete the redundant global variables to ensure that the global variables are unique in the same project. In this case, all users who have the permission to modify the global variables can change the variable values.

Deleting Variables

On the **Global Variables** page, click **Delete** in the **Operation** column of a variable to delete the variable value.

NOTE

- If there are multiple global variables with the same name in the same project under an account, delete the global variables created by the user first. If there are only unique global variables, all users who have the delete permission can delete the global variables.
- After a variable is deleted, the variable cannot be used in SQL statements.

9.2 Permission Management for Global Variables

Scenario

- You can grant permissions on a global variable to users.
- The administrator and the global variable owner have all permissions. You do not need to set permissions for them, and their global variable permissions cannot be modified by other users.
- When setting global variables for a new user, ensure that the region of the user group the user belongs to has the **Tenant Guest** permission. For details

about the Tenant Guest permission and how to apply for the permission, see [Creating a User Group and Assigning Permissions](#) and [System Permissions](#).

Granting Permissions on a Global Variable to a User

1. Choose **Global Configuration > Global Variables**, locate the row containing the desired global variable, and click **Set Permission** in the **Operation** column. On the **User Permissions** page displayed, you can grant permissions for the global variable, set and revoke user permissions.
2. Click **Grant Permission** in the upper right corner of the page.

Figure 9-2 Granting permissions of a global variable to a user

Grant Permission

* Username

Select the permissions to be granted to the user

Select all

Update Delete Grant Permission

Revoke Permission View Other User's Permissions

Table 9-2 Global variable parameters

Parameter	Description
Username	Name of the IAM user to whom permissions on the datasource connection are to be granted. NOTE The username is the name of an existing IAM user.
Permissions to be granted to the user	<ul style="list-style-type: none"> • Update: This permission allows you to update the global variable. • Delete: This permission allows you to delete the global variable. • Grant Permission: This permission allows you to grant permissions of the global variable to other users. • Revoke Permission: This permission allows you to revoke the global variable permissions that other users have but cannot revoke the global variable owner's permissions. • View Other User's Permissions: This permission allows you to view the global variable permissions of other users.

Granting Global Variable Permissions

Click **Set Permission** in the **Operation** column of the sub-user to modify their permissions. [Table 9-2](#) lists the permission parameters.

If all permission options are grayed out, you are not allowed to change permissions on this global variable. You can apply to the administrator, group owner, or other authorized users for required permissions on the global variable.

Revoking Global Variable Permissions

Click **Revoke Permission** in the **Operation** column of a sub-user to revoke their permissions. After this operation, the sub-user does not have any permission on the global variable.

9.3 Project Configurations

Scenario

When you use a DLI job, you need an OBS bucket to store temporary data, such as job logs and results generated during job running.

Configure a DLI job bucket on the **Global Configuration > Project** page.

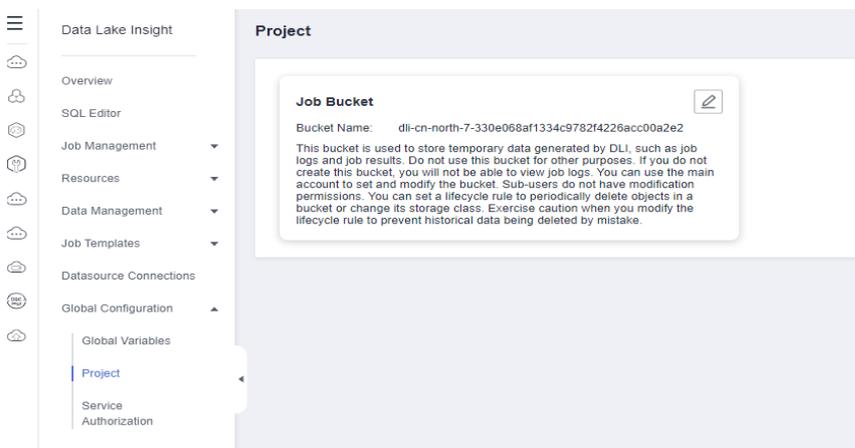
Notes

- The OBS bucket must be used to store temporary data generated by DLI, such as job logs and job results.
- Do not use the OBS bucket for other purposes.
- The OBS bucket must be set and modified by the main account. Member users do not have the permission.
- If the bucket is not configured, you will not be able to view job logs.
- You can create lifecycle rules to automatically delete objects or change storage classes for objects that meet specified conditions.
- Inappropriate modifications of the job bucket may lead to loss of historical data.

Procedure

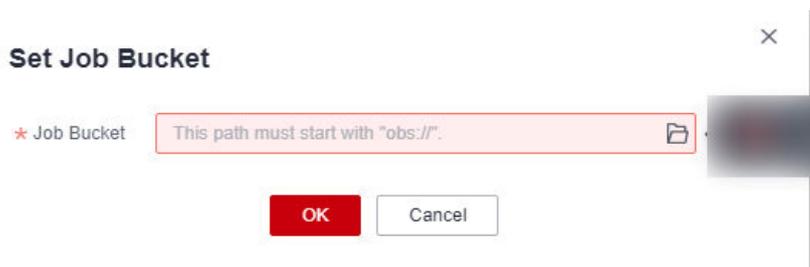
1. In the navigation pane of the DLI console, choose **Global Configuration > Project**.
2. On the **Project** page, click  next to **Job Bucket** to configure bucket information.

Figure 9-3 Project



3. Click  to view available buckets.
4. Select the bucket for storing the temporary data of the DLI job and click **OK**. Temporary data generated during DLI job running will be stored in the OBS bucket.

Figure 9-4 Setting the job bucket



9.4 Service Authorization

Prerequisites

Only the tenant account or a subaccount of user group **admin** can authorize access.

Procedure

After entering the DLI management console, you are advised to set agency permissions to ensure that DLI can be used properly.

If you need to adjust the agency permissions, modify them on the **Service Authorization** page. For details about the required agency permissions, see [Table 9-3](#).

1. Select required agency permissions and click **Update Authorization**. Only the tenant account or a subaccount of user group **admin** can authorize access. If the message "Agency permissions updated" is displayed, the update is successful.

- Once service authorization has succeeded, an agency named **dli_admin_agency** on IAM will be created. Go to the agency list to view the details. Do not delete **dli_admin_agency**.

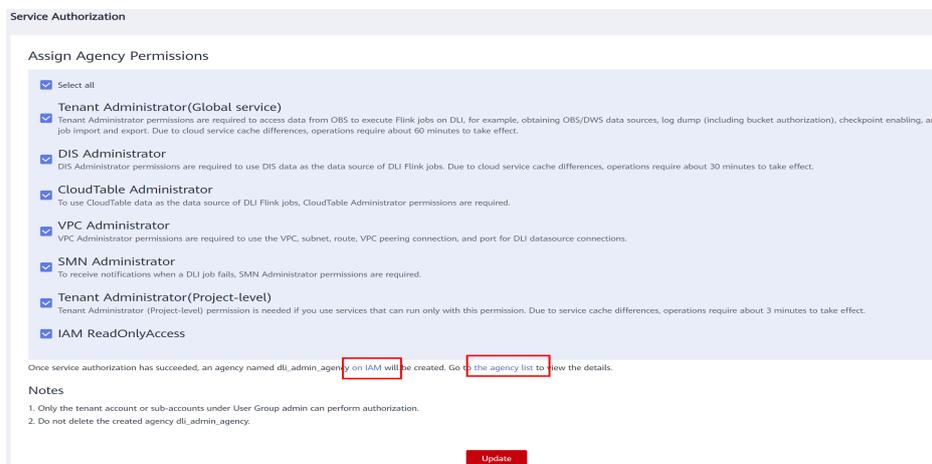


Table 9-3 DLI agency permissions

Permission	Details	Remarks
Tenant Administrator (global service)	Tenant Administrator permissions are required to access data from OBS to execute Flink jobs on DLI, for example, obtaining OBS/DWS data sources, log dump (including bucket authorization), checkpointing enabling, and job import and export.	Due to cloud service cache differences, permission setting operations require about 60 minutes to take effect.
DIS Administrator	DIS Administrator permissions are required to use DIS data as the data source of DLI Flink jobs.	Due to cloud service cache differences, permission setting operations require about 30 minutes to take effect.
CloudTable Administrator	To use CloudTable data as the data source of DLI Flink jobs, CloudTable Administrator permissions are required.	Due to cloud service cache differences, permission setting operations require about 3 minutes to take effect.
VPC Administrator	VPC Administrator permissions are required to use the VPC, subnet, route, VPC peering connection, and port for DLI datasource connections.	Due to cloud service cache differences, permission setting operations require about 3 minutes to take effect.

Permission	Details	Remarks
SMN Administrator	To receive notifications when a DLI job fails, SMN Administrator permissions are required.	Due to cloud service cache differences, permission setting operations require about 3 minutes to take effect.
IAM ReadOnlyAccess	To authorize IAM users who have not logged in to DLI, you need to obtain their information. Therefore, the IAM ReadOnlyAccess permission is required.	To authorize IAM users who have not logged in to DLI, you need to obtain their information. So, the IAM ReadOnlyAccess permission is required.

10 Custom Images

10.1 Overview

Introduction

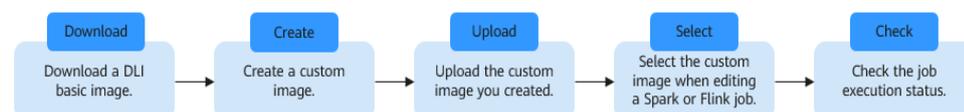
DLI supports clusters deployed in containers. In a container cluster, components related to Spark jobs and Flink jobs run in containers. You can download custom images provided by DLI to change the container running environments of Spark jobs and Flink jobs. For example, you can add a Python package or C library related to machine learning to a custom image to help you extend functions.

Use Restrictions

- The DLI container queue must be used and it must be a dedicated queue.
- The basic image provided by DLI must be used.
- You cannot modify the DLI components and directories in the basic image.
- Only Spark jobs and Flink Jar jobs are supported.

Use Process

Figure 10-1 Process of using a custom image



10.2 Downloading a Basic Image

DLI offers various versions of basic images.

Table 10-1 Downloading a basic image

Version	Download Address
Spark 3.1.1	spark-3.1.1-1.0.0-x86_64.tar.gz
Spark 2.4.5	spark-2.4.5-1.0.1-x86_64.tar.gz
Flink 1.12	flink-1.12.2-1.0.0-x86_64.tar.gz

10.3 Using Custom Images

Scenario

You can use DLI's Spark or Flink basic images and a Dockerfile to package the dependencies (files, JAR files, or software) required to run jobs into a custom image. You can then publish the image to SWR and select it in the DLI job editing page to run jobs. You can embed private capabilities into custom images to enhance the functions and performance of jobs.

NOTE

To use the custom image function, you need to have basic knowledge of [Docker](#).

Restrictions

- Only containerized queues are supported.
- Only JAR jobs (Spark and Flink Jar jobs) are supported.

Procedure for Creating a Custom Image

The following describes how to package TensorFlow into an image to generate a custom image with TensorFlow installed. Then, you can use the image to run jobs in DLI.

Step 1 Prepare the container environment.

For details, see "Step 1: Install the Container Engine" in [Uploading an Image Through a Container Engine Client](#).

Step 2 Log in to the prepared container environment as user **root** and run a command to obtain the basic image.

This section uses a Spark basic image as an example to describe how to download it to the prepared container environment using Wget.

wget *Image download address in [Downloading a Basic Image](#)*

For example, to download the basic image of Spark 2.4.5, run the following command:

```
wget https://uquery-image.obs.xxx.com:443/1.0.0/spark-2.4.5-1.0.0-x86_64.tar.gz
```

Step 3 Load a custom image from local storage.

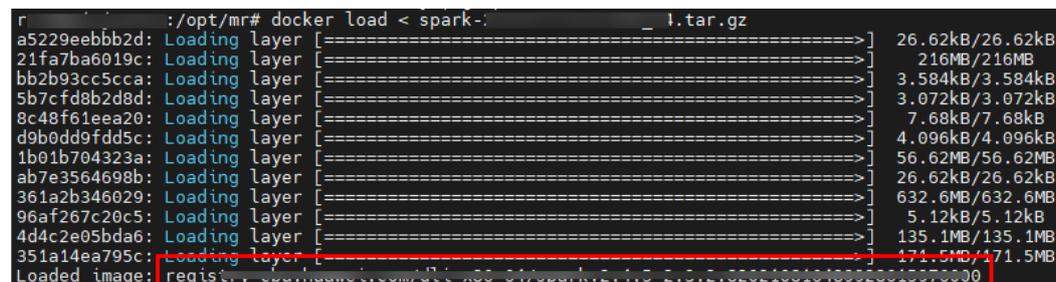
docker load < *Name of the basic image package you have downloaded*

For example, run the following command to load the Spark 2.4.5 basic image:

```
docker load < spark-2.4.5-1.0.0-x86_64.tar.gz
```

Note the basic image loading address following **Loaded image**: as shown in [Figure 10-2](#). This address will be used in subsequent steps.

Figure 10-2 Image loading address



```
r      :/opt/mr# docker load < spark-1.tar.gz
a5229e8bb2d: Loading layer [=====>] 26.62kB/26.62kB
21fa7ba6019c: Loading layer [=====>] 216MB/216MB
bb2b93cc5cca: Loading layer [=====>] 3.584kB/3.584kB
5b7cfd8b2d8d: Loading layer [=====>] 3.072kB/3.072kB
8c48f61eea20: Loading layer [=====>] 7.68kB/7.68kB
d9b0dd9fdd5c: Loading layer [=====>] 4.096kB/4.096kB
1b01b704323a: Loading layer [=====>] 56.62MB/56.62MB
ab7e3564698b: Loading layer [=====>] 26.62kB/26.62kB
361a2b346029: Loading layer [=====>] 632.6MB/632.6MB
96af267c20c5: Loading layer [=====>] 5.12kB/5.12kB
4d4c2e05bda6: Loading layer [=====>] 135.1MB/135.1MB
351a14ea795c: Loading layer [=====>] 171.5MB/171.5MB
Loaded image: registry-cbu.huawei.com/dli-x86_64/spark-2.4.5-2.3.3-sw-xxx/testdli0617/spark:2.4.5.test
```

Step 4 Access SWR.

1. Log in to the SWR console.
2. In the navigation pane on the left, choose **Dashboard** and click **Generate Login Command** in the upper right corner. On the displayed page, click  to copy the login command.
3. Run the login command on the VM where the container engine is installed.

Step 5 Create an organization. If an organization has been created, skip this step.

1. Log in to the SWR console.
2. In the navigation pane on the left, choose **Organization Management**. On the displayed page, click **Create Organization** in the upper right corner.
3. Enter the organization name and click **OK**.

Step 6 Add a tag to the loaded basic image.

docker tag [Step 3](#) *Image loading address* [Image repository address]/
[Organization name]/[Image name:Version]

- [Image repository address]: The domain name at the end of the login command in [Step 4](#) is the image repository address, which can be obtained on the SWR console.
- [Organization name]: Name of the organization created in [Step 5](#)
- [Image name:Version]: Customized name and version of the image.

The following is an example:

```
docker tag registry-cbu.huawei.com/dli-x86_64/spark:2.4.5-2.3.3 sw-xxx/testdli0617/spark:2.4.5.test
```

Step 7 Upload the basic image to the image repository.

docker push [Image repository address]/[Organization name]/[Image name:Version]

Set [Image repository address]/[Organization name]/[Image name:Version] the same as those in [Step 6](#).

The following is an example:

```
docker push swr.xxx/testdli0617/spark:2.4.5.test
```

Step 8 Write a Dockerfile.

vi Dockerfile

Pack TensorFlow into the image as follows:

```
ARG BASE_IMG=spark:latest

FROM ${BASE_IMG} as builder
USER root
RUN set -ex && \
    mkdir -p /root/pip && \ //Configure the available repository address of pip.
    echo "[global]" >> /root/.pip/pip.conf && \ //Configure the Huawei image source.
    echo "index-url = https://mirrors.huaweicloud.com/repository/pypi/simple" >> /root/.pip/pip.conf && \ //
    Configure the Huawei image source.
    RUN pip3 install --user --no-cache-dir tensorflow==1.13.1

FROM ${BASE_IMG}
USER root
COPY --from=builder --chown=omm:ficommon /root/.local /home/omm/.local //Copy the content in the
temporary image builder that incorporates the algorithm package to the basic image.
USER omm
```

The following steps are included:

1. Set the available repository address of pip.
2. Use pip3 to install the TensorFlow algorithm package.
3. Copy the content in the temporary image builder where the algorithm package is installed to the basic image (this step is to reduce the image size) to generate the final custom image.

Step 9 Use Dockerfile to generate a custom image.

Format of the image packaging command:

```
docker build -t [Custom organization name]/[Custom image name]:[Image version] --build-arg
BASE_IMG=[DLI basic image path] -f Dockerfile
```

The basic image path is [*Image repository address*]/[*Organization name*]/[*Image name:Version*] in [Step 6](#).

The following is an example:

```
docker build -t mydli/spark:2.4 --build-arg BASE_IMG=swr.xxx/testdli0617/spark:2.4.5.test -f Dockerfile .
```

Step 10 Add a tag to the custom image.

docker tag [*Organization name*]/[*Image name*]:[*Image version*][*Image repository address*]/[*Organization name*]/[*Image name:version*] in [Step 9](#)

The following is an example:

```
docker tag mydli/spark:2.4 swr.xxx/testdli0617/spark:2.4.5.tensorflow
```

Step 11 Upload the custom image.

docker push [*Image repository address*]/[*Organization name*]/[*Image name:Version*]

Set [*Image repository address*]/[*Organization name*]/[*Image name:Version*] the same as those in [Step 10](#).

The following is an example:

```
docker push swr.xxx/testdli0617/spark:2.4.5.tensorflow
```

Step 12 When submitting a Spark or Flink JAR job in DLI, select a custom image.

- Open the Spark job or Flink job editing page on the management console, select the uploaded and shared image from the custom image list, and run the job.

If you select a non- shared image, the system displays a message indicating that the image is not authorized. You can use the image only after it is authorized. Click **Authorize** as prompted as shown in **Figure 10-4**. Configure other job execution parameters and execute the job.

Figure 10-3 Selecting a custom image on the DLI Spark job editing page

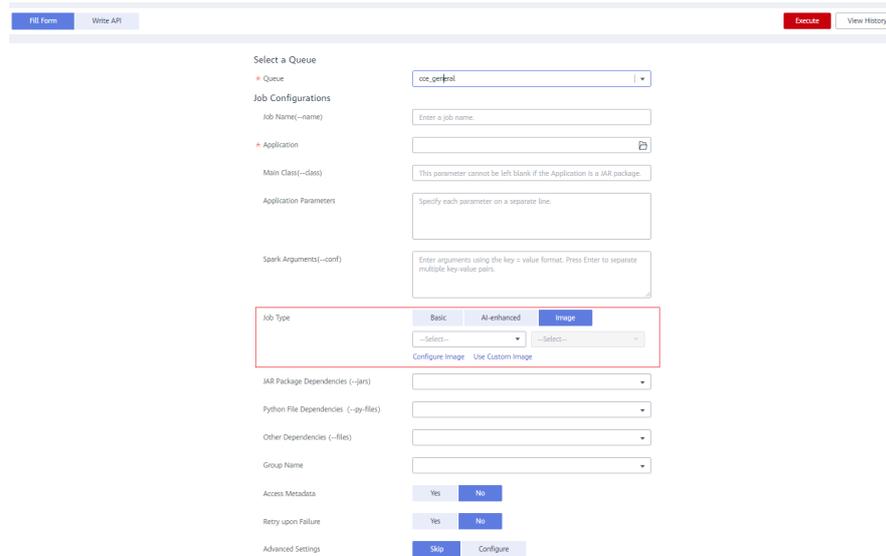


Figure 10-4 Authorizing a Spark job image

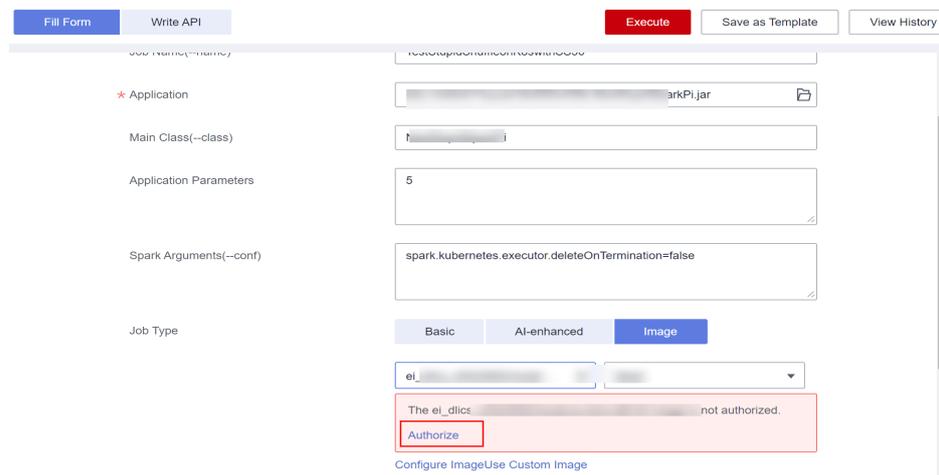


Figure 10-5 Selecting a custom image on the DLI Flink Jar job editing page

test_s [redacted]

ID: 93229 | Job Type: Flink Jar

* Queue

* Application [View Built-in Dependencies](#)

Main Class

Default main class is specified by the Manifest file of the application.

Class Arguments

JAR Package Dependenc... [View Built-in Dependencies](#)

Other Dependencies

Job Type

[Configure ImageUse Custom Image](#)

Flink Version

Runtime Configuration

- Specify the image parameter in job parameters on API to use a custom image to run a job.

For details about Spark jobs, see [Creating a Batch Processing Job](#). For details about Flink Jar jobs, see [Creating a Flink Jar Job](#).

----End

10.4 Basic Image Components

- Basic image system version:
euleros 2.9.3
- Components of the basic image system:

Table 10-2 Basic image system components

System Component	Edition
python3	3.7.4
pip3	20.1.1
java	1.8.0_242
glibc	2.28-36.oe1

- DLI basic image component list:

Table 10-3 DLI basic image components

Component	DLI Spark Basic Image	DLI Spark AI basic image
spark	2.4.5	2.4.5
numpy	1.16.1	1.16.1
pandas	0.24.2	0.24.2
pyarrow	0.11.1	0.11.1

11 Permissions Management

11.1 Overview

DLI has a comprehensive permission control mechanism and supports fine-grained authentication through Identity and Access Management (IAM). You can create policies in IAM to manage DLI permissions. You can use both the DLI's permission control mechanism and the IAM service for permission management.

Application Scenarios of IAM Authentication

When using DLI on the public cloud, enterprise users need to manage DLI resources (queues) used by employees in different departments, including creating, deleting, using, and isolating resources. In addition, data of different departments needs to be managed, including data isolation and sharing.

DLI uses IAM for refined enterprise-level multi-tenant management. IAM provides identity authentication, permissions management, and access control, helping you securely access to your public cloud resources.

With IAM, you can use your public cloud account to create IAM users for your employees, and assign permissions to the users to control their access to specific resource types. For example, some software developers in your enterprise need to use DLI resources but must not delete them or perform any high-risk operations. To achieve this result, you can create IAM users for the software developers and grant them only the permissions required for using DLI resources.

NOTE

For a new user, you need to log in for the system to record the metadata before using DLI.

IAM is free of charge. You pay only for the resources you use.

If your public cloud account does not need individual IAM users for permissions management, skip this chapter.

DLI System Permissions

[Table 11-1](#) lists all the system-defined roles and policies supported by DLI.

Type: There are roles and policies.

- **Roles:** A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. Only a limited number of service-level roles are available. When using roles to grant permissions, you also need to assign other roles on which the permissions depend. However, roles are not an ideal choice for fine-grained authorization and secure access control.
- **Policies:** A type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, you can grant DLI users only the permissions for managing a certain type of ECSs.

For details about the system policies you need to perform common SQL operations, see [Common Operations Supported by DLI System Policy](#).

Table 11-1 DLI system permissions

Role/Policy Name	Description	Category
DLI FullAccess	Full permissions for DLI.	System-defined policy
DLI ReadOnlyAccess	Read-only permissions for DLI. With read-only permissions, you can use DLI resources and perform operations that do not require fine-grained permissions. For example, create global variables, create packages and package groups, submit jobs to the default queue, create tables in the default database, create datasource connections, and delete datasource connections.	System-defined policy
Tenant Administrator	Tenant administrator <ul style="list-style-type: none">• Administer permissions for managing and accessing all cloud services. After a database or a queue is created, the user can use the ACL to assign rights to other users.• Scope: project-level service	System-defined role
DLI Service Admin	DLI administrator <ul style="list-style-type: none">• Administer permissions for managing and accessing the queues and data of DLI. After a database or a queue is created, the user can use the ACL to assign rights to other users.• Scope: project-level service	System-defined role

For details, see [Creating an IAM User and Granting Permissions, How Do I Create an IAM user?](#) and [How Do I Modify a User Policy?](#)

DLI Permission Types

Table 11-2 lists the DLI service permissions. For details about the resources that can be controlled by DLI, see [Table 11-7](#).

Table 11-2 DLI permission types

Permission Type	Subtype	Console Operations	SQL Syntax	API Definition
Queue Permissions	Queue management permissions	For details, see Queue Permission Management .	None	For details, see "Granting Users with the Queue Usage Permission" in the <i>Data Lake Insight API Reference</i> .
	Queue usage permission			
Data Permissions	Database permissions	For details, see Managing Database Permissions and Managing Table Permissions .	For details, see SQL Syntax of Batch Jobs > Data Permissions Management > Data Permissions List in the <i>Data Lake Insight SQL Syntax Reference</i> .	For details, see Permission-related APIs > Granting Users with the Data Usage Permission in the <i>Data Lake Insight API Reference</i> .
	Table permissions			
	Column permissions			
Job Permissions	Flink job permissions	For details, see Managing Flink Job Permissions .	None	For details, see Permission-related APIs > Granting Users with the Data Usage Permission in the <i>Data Lake Insight API Reference</i> .

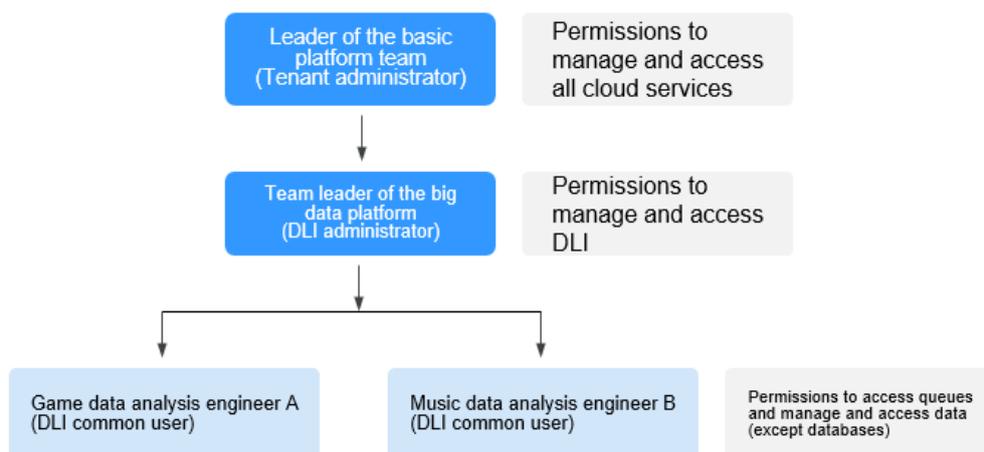
Permission Type	Subtype	Console Operations	SQL Syntax	API Definition
Package Permissions	Package group permissions	For details, see Managing Permissions on Packages and Package Groups .	None	For details, see Permission-related APIs > Granting Users with the Data Usage Permission in the <i>Data Lake Insight API Reference</i> .
	Package permissions			
Datasource Connection Permissions	Datasource connection permissions	For details, see Datasource Authentication on Permission Management .	None	For details, see Permission-related APIs > Granting Users with the Data Usage Permission in the <i>Data Lake Insight API Reference</i> .

Examples

An Internet company mainly provides game and music services. DLI is used to analyze user behaviors and assist decision making.

As shown in [Figure 11-1](#), the **Leader of the Basic Platform Team** has applied for a **Tenant Administrator** account to manage and use public cloud services. Since the **Big Data Platform Team** needs DLI for data analysis, the **Leader of the Basic Platform Team** adds a subaccount with the permission of **DLI Service Admin** to manage and use DLI. The **Leader of the Basic Platform Team** creates a **Queue A** and assigns it to **Data Engineer A** to analyze the gaming data. A **Queue B** is also assigned to **Data Engineer B** to analyze the music data. Besides granting the queue usage permission, the **Leader of the Basic Platform Team** grants data (except the database) management and usage permissions to the two engineers.

Figure 11-1 Granting permissions



The **Data Engineer A** creates a table named **gameTable** for storing game prop data and a table named **userTable** for storing game user data. The music service is a new service. To explore potential music users among existing game users, the **Data Engineer A** assigns the query permission on the **userTable** to the **Data Engineer B**. In addition, **Data Engineer B** creates a table named **musicTable** for storing music copyrights information.

Table 11-3 describes the queue and data permissions of **Data Engineer A** and **Data Engineer B**.

Table 11-3 Permission description

User	Data Engineer A (game data analysis)	Data Engineer B (music data analysis)
Queues	Queue A (queue usage permission)	Queue B (queue usage permission)
Data (Table)	gameTable (table management and usage permission)	musicTable (table management and usage permission)
	userTable (table management and usage permission)	userTable (table query permission)

NOTE

The queue usage permission includes job submitting and terminating permissions.

11.2 Creating an IAM User and Granting Permissions

You can use Identity and Access Management (IAM) to implement fine-grained permissions control on DLI resources. For details, see [Overview](#).

If your cloud account does not need individual IAM users, then you may skip over this chapter.

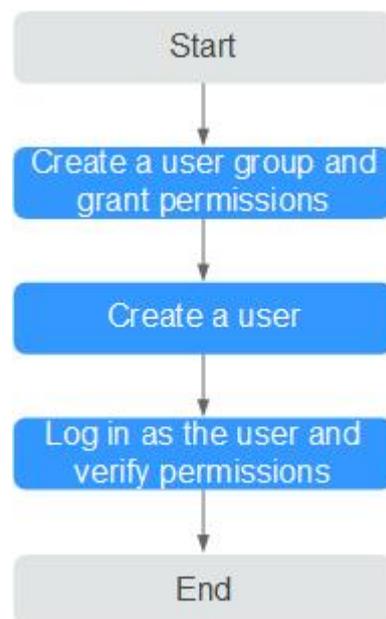
This section describes how to create an IAM user and grant DLI permissions to the user. [Figure 11-2](#) shows the procedure.

Prerequisites

Before assigning permissions to user groups, you should learn about system policies and select the policies based on service requirements. For details about system permissions supported by DLI, see [DLI System Permissions](#).

Process Flow

Figure 11-2 Process for granting DLI permissions



1. Create a user group and grant the permission to it.
Create a user group on the IAM console, and assign the **DLI ReadOnlyAccess** permission to the group.
2. Create a user and add the user to the user group.
Create a user on the IAM console and add the user to the group created in [1](#).
3. Log in and verify the permission.
Log in to the management console using the newly created user, and verify that the user's permissions.
 - Choose **Service List > Data Lake Insight**. The DLI management console is displayed. If you can view the queue list on the **Queue Management** page but cannot buy DLI queues by clicking **Buy Queue** in the upper right corner (assume that the current permission contains only **DLI ReadOnlyAccess**), the **DLI ReadOnlyAccess** permission has taken effect.
 - Choose any other service in **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **DLI ReadOnlyAccess** permission has already taken effect.

11.3 Creating a Custom Policy

Custom policies can be created as a supplement to the system policies of DLI. You can add actions to custom policies. For the actions supported for custom policies, see "Permissions Policies and Supported Actions" in the *Elastic Volume Service API Reference*.

You can create custom policies in either of the following two ways:

- Visual editor: Select cloud services, actions, resources, and request conditions without the need to know policy syntax.
- JSON: Create a policy in the JSON format from scratch or based on an existing policy.

. This section describes common DLI custom policies.

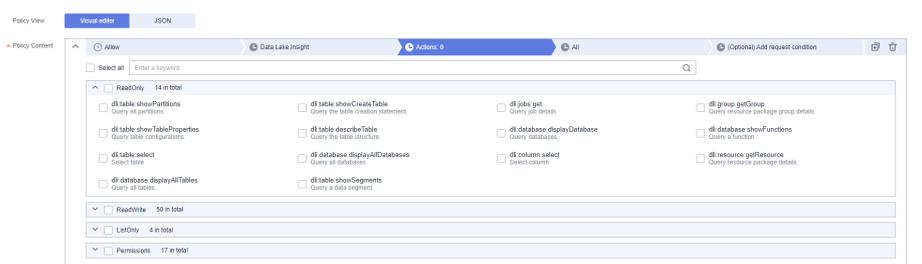
Policy Field Description

The following example assumes that the authorized user has the permission to create tables in all databases in all regions:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dli:database:createTable"
      ],
      "Resource": [
        "dli:*:*:database:*"
      ]
    }
  ]
}
```

- Version
1.1 indicates a fine-grained permission policy that defines permissions required to perform operations on specific cloud resources under certain conditions.
- Effect
The value can be **Allow** and **Deny**. If both **Allow** and **Deny** are found in statements, the **Deny** overrides the **Allow**.
- Action
Specific operation on a resource. A maximum of 100 actions are allowed, as shown in [Figure 11-3](#).

Figure 11-3 DLI actions



 NOTE

- The format is *Service name.Resource type.Action*, for example, **dli:queue:submit_job**.
 - *Service name*: product name, such as **dli**, **evs**, and **vpc**. Only lowercase letters are allowed. Resource types and operations are not case-sensitive. You can use an asterisk (*) to represent all operations.
 - *Resource type*: For details, see [Table 11-7](#).
 - *Action*: action registered in IAM.
- Condition

Conditions determine when a policy takes effect. A condition consists of a condition key and operator.

A condition key is a key in the **Condition** element of a statement. There are global and service-level condition keys.

- Global condition keys (prefixed with **g:**) apply to all actions. For details, see condition key description in Policy Syntax.
- Service-level condition keys apply only to operations of the specific service.

An operator is used together with a condition key to form a complete condition statement. For details, see [Table 11-4](#).

IAM provides a set of DLI predefined condition keys. The following table lists the predefined condition keys of DLI.

Table 11-4 DLI request conditions

Condition Key	Type	Operator	Description
g:CurrentTime	Global	Date and time	Time when an authentication request is received NOTE The time is expressed in the format defined by ISO 8601 , for example, 2012-11-11T23:59:59Z .
g:MFAPresent	Global	Boolean	Whether multi-factor authentication is used during user login
g:UserId	Global	String	ID of the current login user
g:UserName	Global	String	Current login user
g:ProjectName	Global	String	Project that you have logged in to
g:DomainName	Global	String	Domain that you have logged in to
g:ResourceTag	Global	StringEquals	Resource tag value.

- Resource
The format is *Service name:Region:Domain ID:Resource type:Resource path*. The wildcard (*) indicates all options. For details about the resource types and path, see [Table 11-7](#).
Example:
dli:*:*:queue:* indicates all queues.

Creating a Custom Policy

You can set actions and resources of different levels based on scenarios.

1. Define an action.

The format is *Service name:Resource type:Action*. The wildcard is *. Example:

Table 11-5 Action

Action	Description
dli:queue:submit_job	Submission operations on a DLI queue
dli:queue:*	All operations on a DLI queue
dli:*:*	All operations on all DLI resource types

For more information about the relationship between operations and system permissions, see [Common Operations Supported by DLI System Policy](#).

2. Define a resource.

The format is *Service name:Region:Domain ID:Resource type:Resource path*. The wildcard (*) indicates all resources. The five fields can be flexibly set. Different levels of permission control can be set for resource paths based on scenario requirements. If you need to set all resources of the service, you do not need to specify this field. For details about the definition of Resource, see [Table 11-6](#). For details about the resource types and paths in Resource, see [Table 11-7](#).

Table 11-6 Resource

Resource	Description
DLI:*:*:table:databases.dbname.tables.*	DLI, any region, any account ID, all table resources of database dbname
DLI:*:*:database:databases.dbname	DLI, any region, any account ID, resource of database dbname
DLI:*:*:queue:queues.*	DLI, any region, any account ID, any queue resource

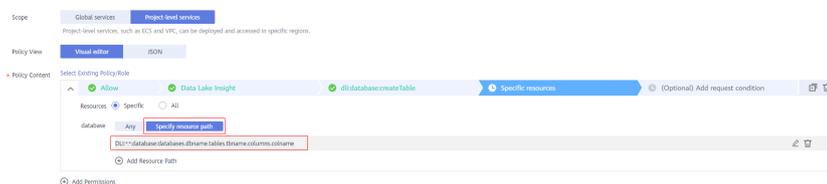
Resource	Description
DLI:*:*:jobs:jobs.flink.1	DLI, any region, any account ID, Flink job whose ID is 1

Table 11-7 DLI resources and their paths

Resource Type	Resource Names	Path
queue	DLI queue	queues.queueName
database	DLI database	databases.dbName
table	DLI table	databases.dbName.tables.tbName
column	DLI column	databases.dbName.tables.tbName.columns.columnName
jobs	DLI Flink job	jobs.flink.jobId
resource	DLI package	resources.resourceName
group	DLI package group	groups.groupName
datasource auth	DLI cross-source authentication information	datasourceauth.name
edsconnections	Enhanced datasource connection	edsconnections. <i>connection ID</i>

- Specific resources:

Figure 11-4 Specific resources



- All resources: all resources of the service

Figure 11-5 All resources



- Combine all the preceding fields into a JSON file to form a complete policy. You can set multiple actions and resources. You can also create a policy on the visualized page provided by IAM. For example:

The authorized user has the permission to create and delete any database, submit jobs for any queue, and delete any table under any account ID in any region of DLI.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dli:database:createDatabase",
        "dli:database:dropDatabase",
        "dli:queue:submitJob",
        "dli:table:dropTable"
      ],
      "Resource": [
        "dli:*:*:database:*",
        "dli:*:*:queue:*",
        "dli:*:*:table:*"
      ]
    }
  ]
}
```

Example Custom Policies

- Example 1: Allow policies
 - Allow users to create tables across all databases in all regions:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dli:database:createTable"
      ],
      "Resource": [
        "dli:*:*:database:*"
      ]
    }
  ]
}
```
 - Allow users to query column **col** in the table **tb** of the database **db** in the region where the user is located:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dli:column:select"
      ],
      "Resource": [
        "dli:xx:*:column:databases.db.tables.tb.columns.col"
      ]
    }
  ]
}
```
- Example 2: Deny policies

A deny policy must be used together with other policies. That is, a user can set a deny policy only after being assigned some operation permissions. Otherwise, the deny policy does not take effect.

If the permissions assigned to a user contain both Allow and Deny actions, the Deny actions take precedence over the Allow actions.

- Deny users to create or delete databases, submit jobs (except the default queue), or delete tables.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "dli:database:createDatabase",
        "dli:database:dropDatabase",
        "dli:queue:submitJob",
        "dli:table:dropTable"
      ],
      "Resource": [
        "dli:*:database:*",
        "dli:*:queue:*",
        "dli:*:table:*"
      ]
    }
  ]
}
```

- Deny users to submit jobs in the demo queue.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "dli:queue:submitJob"
      ],
      "Resource": [
        "dli:*:queue:queues.demo"
      ]
    }
  ]
}
```

- Example 3: Tag authentication. You need to specify an action and bind it to a condition, and specify the key and value of **g:ResourceTag**.

Condition g: ResourceTag indicates a resource with the *key=value* tag. Only the operation on this resource contained in the policy action list is allowed.

The key is case insensitive. Fuzzy match for the value is not supported.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dli:database:dropDatabase",
        "dli:table:select",
        "dli:database:createTable",
        "dli:table:dropTable"
      ],
      "Condition": {
        "StringEquals": {
          "g:ResourceTag/key": [
            "value"
          ]
        }
      }
    }
  ]
}
```

```
}  
  }  
] }  
}
```

11.4 DLI Resources

A resource is an object that exists within a service. You can select DLI resources by specifying their paths.

Table 11-8 DLI resources and their paths

Resource Type	Resource Names	Path
queue	DLI queue	queues.queueName
database	DLI database	databases.dbName
table	DLI table	databases.dbName.tables.tbName
column	DLI column	databases.dbName.tables.tbName.columns.colName
jobs	DLI Flink job	jobs.flink.jobId
resource	DLI package	resources.resourceName
group	DLI package group	groups.groupName
datasourceauth	DLI cross-source authentication information	datasourceauth.name
edsconnections	Enhanced datasource connection	edsconnections. <i>connection ID</i>

11.5 DLI Request Conditions

Request conditions are useful in determining when a custom policy takes effect. A request condition consists of a condition key and operator. Condition keys are either global or service-level and are used in the Condition element of a policy statement. Global condition keys (starting with **g:**) are available for operations of all services, while service-level condition keys (starting with a service name such as **dli**) are available only for operations of a specific service. An operator is used together with a condition key to form a complete condition statement.

IAM provides a set of DLI predefined condition keys. The following table lists the predefined condition keys of *DLI*.

Table 11-9 DLI request conditions

Condition Key	Type	Operator	Description
g:CurrentTime	Global	Date and time	Time when an authentication request is received NOTE The time is expressed in the format defined by ISO 8601 , for example, 2012-11-11T23:59:59Z .
g:MFAPresent	Global	Boolean	Whether multi-factor authentication is used during user login
g:UserId	Global	String	ID of the current login user
g:UserName	Global	String	Current login user
g:ProjectName	Global	String	Project that you have logged in to
g:DomainName	Global	String	Domain that you have logged in to
g:ResourceTag	Global	StringEquals	Resource tag value.

11.6 Common Operations Supported by DLI System Policy

Table 11-10 lists the common operations supported by each system policy of DLI. Choose proper system policies according to this table. For details about the SQL statement permission matrix in DLI in terms of permissions on databases, tables, and roles, see **SQL Syntax of Batch Jobs > Data Permissions Management > Data Permissions List** in the *Data Lake Insight SQL Syntax Reference*.

Table 11-10 Common operations supported by each system policy

Resources	Operation	Description	DLI FullAccess	DLI ReadOnlyAccess	Tenant Administrator	DLI Service Admin
Queue	DROP_QUEUE	Deleting a queue	√	×	√	√
	SUBMIT_JOB	Submitting the job	√	×	√	√

Resources	Operation	Description	DLI FullAccess	DLI ReadOnlyAccess	Tenant Administrator	DLI Service Admin
	CANCEL_JOB	Terminating the job	√	×	√	√
	RESTART	Restarting a queue	√	×	√	√
	GRANT_PRIVILEGE	Granting permissions to the queue	√	×	√	√
	REVOKE_PRIVILEGE	Revoking permissions from the queue	√	×	√	√
	SHOW_PRIVILEGES	Viewing the queue permissions of other users	√	×	√	√
Database	DROP_DATABASE	Deleting a database	√	×	√	√
	CREATE_TABLE	Creating a table	√	×	√	√
	CREATE_VIEW	Creating a view	√	×	√	√
	EXPLAIN	Explaining the SQL statement as an execution plan	√	×	√	√
	CREATE_ROLE	Creating a role	√	×	√	√
	DROP_ROLE	Deleting a role	√	×	√	√
	SHOW_ROLES	Displaying a role	√	×	√	√
	GRANT_ROLE	Binding a role	√	×	√	√
	REVOKE_ROLE	Unbinding the role	√	×	√	√

Resources	Operation	Description	DLI FullAccess	DLI ReadOnlyAccess	Tenant Administrator	DLI Service Admin
	SHOW_USERS	Displaying the binding relationships between all roles and users	√	×	√	√
	GRANT_PRIVILEGE	Granting permissions to the database	√	×	√	√
	REVOKE_PRIVILEGE	Revoking permissions to the database	√	×	√	√
	SHOW_PRIVILEGES	Viewing database permissions of other users	√	×	√	√
	DISPLAY_ALL_TABLES	Displaying tables in the database	√	√	√	√
	DISPLAY_DATABASE	Displaying databases	√	√	√	√
	CREATE_FUNCTION	Creating a function	√	×	√	√
	DROP_FUNCTION	Deleting a function	√	×	√	√
	SHOW_FUNCTIONS	Displaying all functions	√	×	√	√
	DESCRIBE_FUNCTION	Displaying function details	√	×	√	√
Table	DROP_TABLE	Deleting a table	√	×	√	√
	SELECT	Querying a table	√	×	√	√
	INSERT INTO TABLE	Inserting	√	×	√	√

Resources	Operation	Description	DLI FullAccess	DLI ReadOnlyAccess	Tenant Administrator	DLI Service Admin
	ALTER_TABLE_ADD_COLUMNS	Adding a column	√	×	√	√
	INSERT_OVERRWRITE_TABLE	Rewriting	√	×	√	√
	ALTER_TABLE_RENAME	Renaming a table	√	×	√	√
	ALTER_TABLE_ADD_PARTITION	Adding partitions to the partition table	√	×	√	√
	ALTER_TABLE_RENAME_PARTITION	Renaming a table partition	√	×	√	√
	ALTER_TABLE_DROP_PARTITION	Deleting partitions from a partition table	√	×	√	√
	SHOW_PARTITIONS	Displaying all partitions	√	×	√	√
	ALTER_TABLE_RECOVER_PARTITION	Restoring table partitions	√	×	√	√
	ALTER_TABLE_SET_LOCATION	Setting the partition path	√	×	√	√
	GRANT_PRIVILEGE	Granting permissions to the table	√	×	√	√
	REVOKE_PRIVILEGE	Revoking permissions from the table	√	×	√	√
	SHOW_PRIVILEGES	Viewing table permissions of other users	√	×	√	√

Resources	Operation	Description	DLI FullAccess	DLI ReadOnlyAccess	Tenant Administrator	DLI Service Admin
	DISPLAY_TABLE	Displaying a table	√	√	√	√
	DESCRIBE_TABLE	Displaying table information	√	×	√	√

12 Other Common Operations

12.1 Importing Data to a DLI Table

Importing Data Using OBS

On the DLI management console, you can import data stored on OBS to DLI tables from **Data Management > Databases and Tables > Table Management** and **SQL Editor** pages. For details, see [Importing Data to the Table](#).

Importing Data Using CDM

Use the Cloud Data Migration (CDM) service to import data from OBS to DLI. You need to create a CDM queue first.

Pay attention to the following configurations:

- The VPC to which the DLI account belongs is the same as the VPC of the CDM queue.
- You need to create two links, including a DLI link and an OBS link.
- The format of the file to be transmitted can be **CSV** or **JSON**.

12.2 Viewing Monitoring Metrics

Description

This section describes metrics reported by DLI to Cloud Eye as well as their namespaces and dimensions. You can use the management console or APIs provided by Cloud Eye to query the metrics of the monitored object and alarms generated for DLI.

Namespace

SYS.DLI

Metric

Table 12-1 DLI metrics

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
queue_cu_num	CU usage of a queue	Displays the number of CUs applied by the user queue	≥ 0	Queues	5 minutes
queue_job_launching_num	Number of Jobs Being Submitted	Displays the number of jobs in the Submitting state in the user queue.	≥ 0	Queues	5 minutes
queue_job_running_num	Number of Running Jobs	Displays the number of running jobs in the user queue.	≥ 0	Queues	5 minutes
queue_job_succeed_num	Number of Finished Jobs	Displays the number of completed jobs in the user queue.	≥ 0	Queues	5 minutes
queue_job_failed_num	Failed Jobs	Displays the number of failed jobs in the user queue.	≥ 0	Queues	5 minutes

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
queue_job_cancelled_num	Number of Canceled Jobs	Displays the number of canceled jobs in the user queue.	≥ 0	Queues	5 minutes
queue_cpu_usage	Queue CPU Usage	Displays the CPU usage of user queues.	0-100	Queues	5 minutes
queue_disk_usage	Queue Disk Usage	Displays the disk usage of user queues.	0-100	Queues	5 minutes
queue_disk_used	Max Disk Usage	Displays the maximum disk usage of user queues.	0~100	Queues	5 minutes
queue_mem_usage	Queue Memory Usage	Displays the memory usage of user queues.	0-100	Queues	5 minutes
queue_mem_used	Used Memory	Displays the memory usage rate of the user queues.	≥ 0	Queues	5 minutes
flink_read_records_per_second	Flink Job Data Read Rate	Displays the data input rate of a Flink job for monitoring and debugging.	≥ 0	Flink jobs	10 seconds

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
flink_write_records_per_second	Flink Job Data Write Rate	Displays the data output rate of a Flink job for monitoring and debugging.	≥ 0	Flink jobs	10 seconds
flink_read_records_total	Flink Job Total Data Read	Displays the total number of data inputs of a Flink job for monitoring and debugging.	≥ 0	Flink jobs	10 seconds
flink_write_records_total	Flink Job Total Data Write	Displays the total number of output data records of a Flink job for monitoring and debugging.	≥ 0	Flink jobs	10 seconds
flink_read_bytes_per_second	Flink Job Byte Read Rate	Displays the number of input bytes per second of a Flink job.	≥ 0	Flink jobs	10 seconds
flink_write_bytes_per_second	Flink Job Byte Write Rate	Displays the number of output bytes per second of a Flink job.	≥ 0	Flink jobs	10 seconds

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
flink_read_bytes_total	Flink Job Total Read Byte	Displays the total number of input bytes of a Flink job.	≥ 0	Flink jobs	10 seconds
flink_write_bytes_total	Flink Job Total Write Byte	Displays the total number of output bytes of a Flink job.	≥ 0	Flink jobs	10 seconds
flink_cpu_usage	Flink Job CPU Usage	Displays the CPU usage of Flink jobs.	0-100	Flink jobs	10 seconds
flink_mem_usage	Flink Job Memory Usage	Displays the memory usage of Flink jobs.	0-100	Flink jobs	10 seconds
flink_max_operator_latency	Flink Job Max Operator Latency	Displays the maximum operator delay of a Flink job. The unit is ms .	≥ 0	Flink jobs	10 seconds

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
flink_max_operator_backpressure_level	Flink Job Maximum Operator Backpressure	Displays the maximum operator backpressure value of a Flink job. A larger value indicates severer backpressure. 0: OK 50: low 100: high	0-100	Flink jobs	10 seconds

Dimension

Table 12-2 Dimension

Key	Value
queue_id	Queue
flink_job_id	Flink job

Viewing DLI Monitoring Metrics on Cloud Eye

1. Search for Cloud Eye on the management console.
2. In the navigation pane on the left of the Cloud Eye console, click **Cloud Service Monitoring > Data Lake Insight**.
3. Select a queue to view its information.

12.3 DLI Operations That Can Be Recorded by CTS

With CTS, you can record operations associated with DLI for later query, audit, and backtrack operations.

Table 12-3 DLI operations that can be recorded by CTS

Operation	Resource Type	Trace Name
Creating a database	database	createDatabase
Deleting a database	database	deleteDatabase
Modifying the Database Owner	database	alterDatabaseOwner
Creating a table	table	createTable
Deleting tables	table	deleteTable
Exporting table data	table	exportData
Importing table data	table	importData
Modifying the owner of a table	table	alterTableOwner
Creating a queue	queue	createQueue
Deleting a queue	queue	dropQueue
Granting permissions to a queue	queue	shareQueue
Modifying a Queue CIDR Block	queue	replaceQueue
Restarting a queue	queue	queueActions
Scaling out/in a queue	queue	queueActions
Submitting a job (SQL)	queue	submitJob
Canceling a job (SQL)	queue	cancelJob
Granting DLI the permission to access OBS buckets	obs	obsAuthorize
Checking the SQL syntax	job	checkSQL
Creating a job	job	createJob
Updating a job	job	updateJob
Deleting a job	job	deleteJob
Purchasing CUH packages	order	orderPackage
Freezing resources	resource	freezeResource
Unfreezing resources	resource	unfreezeResource
Terminating resources	resource	deleteResource
Clearing resources	resource	cleanResource

Operation	Resource Type	Trace Name
Granting data permissions	data	dataAuthorize
Granting permissions on other projects	data	authorizeProjectData
Exporting query results	data	storeJobResult
Saving a SQL template	sqlTemplate	saveSQLTemplate
Updating a SQL template	sqlTemplate	updateSQLTemplate
Deleting a SQL template	sqlTemplate	deleteSQLTemplate
Creating a Flink template	flinkTemplate	createStreamTemplate
Updating a Flink template	flinkTemplate	createStreamTemplate
Deleting a Flink template	flinkTemplate	deleteStreamTemplate
Creating a data upload task	uploader	createUploadJob
Obtaining the authentication to perform a data upload task	uploader	getUploadAuthInfo
Submitting a data upload task	uploader	commitUploadJob
Creating a datasource authentication and uploading a certificate	authInfo	uploadAuthInfo
Updating a datasource authentication	authInfo	updateAuthInfo
Deleting a datasource authentication	authInfo	deleteAuthInfo
Updating the quota	quota	updateQuota
Uploading a resource package	pkgResource	uploadResources
Deleting a resource package	pkgResource	deleteResource
Creating a basic datasource connection	datasource	createDatasourceConn
Deleting a basic datasource connection	datasource	deleteDatasourceConn
Reactivating a basic datasource connection	datasource	reactivateDSConnection

Operation	Resource Type	Trace Name
Creating an enhanced datasource connection	datasource	createConnection
Deleting an enhanced datasource connection	datasource	getConnection
Binding a queue	datasource	associateQueueToDataSourceConn
Unbinding a queue	datasource	disassociateQueueToDataSourceConn
Modifying the host information	datasource	updateHostInfo
Adding a route	datasource	addRoute
Deleting a route	datasource	deleteRoute
Creating a topic	smn	createTopic
Creating an agency	agency	createAgencyV2
Creating a batch processing job	batch	createBatch
Canceling a batch processing job	batch	cancelBatch
Creating a session	session	createSession
Deleting a session	session	deleteSession
Creating a statement	statement	createStatement
Canceling execution of a statement	statement	cancelStatement
Creating a global variable	globalVar	createGlobalVariable
Deleting a global variable	globalVar	deleteGlobalVariable
Modifying a global variable	globalVar	updateGlobalVariable

For details about how to enable CTS and view trace details, see the [Cloud Trace Service Getting Started](#).

For details about key fields in the CTS trace structure, see [Trace Structure](#) in the *Cloud Trace Service User Guide*.

12.4 Quotas

What Is a Quota?

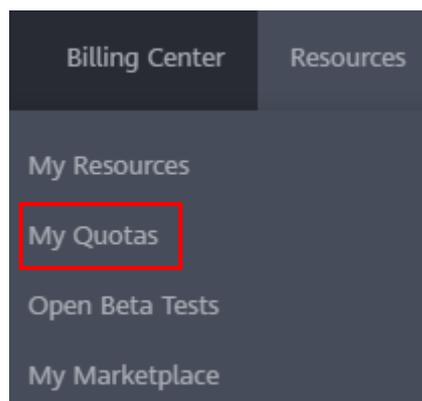
A quota limits the quantity of a resource available to users, thereby preventing spikes in the usage of the resource.

You can also request for an increased quota if your existing quota cannot meet your service requirements.

How Do I View My Quotas?

1. Log in to the management console.
2. Click  in the upper left corner and select **Region** and **Project**.
3. In the upper right corner of the page, choose **Resources** > **My Quotas**.
The **Service Quota** page is displayed.

Figure 12-1 My quotas



4. View the used and total quota of each type of resources on the displayed page.
If a quota cannot meet service requirements, increase a quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources** > **My Quotas**.
The **Service Quota** page is displayed.

Figure 12-2 My quotas



3. Click **Increase Quota**.
4. On the **Create Service Ticket** page, configure parameters as required. In the **Problem Description** area, fill in the content and reason for adjustment.
5. Select the agreement and click **Submit**.

A Change History

Released On	What's New
2023-12-11	Modified the following section: <ul style="list-style-type: none">Removed information about sensitive variables from Global Variables.
2023-11-07	Optimized the following sections: <ul style="list-style-type: none">Modified the introduction to datasource authentication and recommended this authentication in Introduction.Added the description that datasource authentication is recommended for storing cross-source credentials in cross-source scenarios to Overview.
2023-08-28	Added the following section: <ul style="list-style-type: none">Allocating a Queue to an Enterprise Project
2023-08-02	Modified the following content: <ul style="list-style-type: none">Modified the description of sensitive variables in "Global Variables" in Global Configuration.Added Permission Management for Global Variables. Modified the following sections: <ul style="list-style-type: none">Modified the description of the SQL job editing window in SQL Editor.Modified the content of the job management parameter table in SQL Job Management.
2023-06-30	Optimized the following sections: <ul style="list-style-type: none">Enhanced Datasource ConnectionsDatasource Authentication

Released On	What's New
2023-04-17	Modified the following content: <ul style="list-style-type: none">Modified the note about the tag key, tag value, and tag in Managing Queue Tags.
2023-03-03	Added table: Table 3-19
2023-02-23	Modified the following content: <ul style="list-style-type: none">Deleted the content related to Flink job debugging.Performing Operations on a Flink Job: Modified the description of the Auto Restart upon Exception parameter in the runtime configuration. Flink SQL jobs and Flink OpenSource SQL jobs cannot be restored from checkpoints.
2022-09-15	This issue is the first official release.