


Stream Compute Service

Releases Notes and Announcements

Product Documentation



Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Releases Notes and Announcements

Release Notes

Releases Notes and Announcements

Release Notes

Last updated : 2024-06-27 11:53:37

April 2023

| Update | Description | Release Date | Documentation |
|--|---|--------------|--|
| Enhanced Flink UI security control | Flink UI supports access allowlist configuration. | 2023-04-17 | Viewing the Flink UI of a Job |
| Optimized the ClickHouse connector | Optimization is made to alleviate this issue: Joining of multiple streams causes many data changes. As a result, the data sink has much pressure on ClickHouse. | 2023-04-17 | ClickHouse Data Warehouse |
| Optimized job log level configurations | Users can configure the log level of jobs in batches. | 2023-04-17 | Configuring Running Log Collection for a Job |
| Optimized database connection via JDBC | The error in splicing Oracle JDBC statements using composite primary keys is fixed. | 2023-04-17 | JDBC Connector |
| Optimized capabilities related to monitoring metrics | The system supports display and alarm configuration for Kafka job latency metrics. | 2023-04-17 | Monitoring Metric List |
| Optimized FileSystem | The system supports compression formats such as Snappy for sinking data in Raw format to FileSystem. | 2023-04-17 | FileSystem |

March 2023

| Update | Description | Release Date | Documentation |
|-----------|---|--------------|--------------------------|
| Optimized | The logic for deleting dependencies and | 2023-03- | Managing |

| | | | |
|---|--|------------|--------------------------------------|
| dependency management | dependency versions is optimized. | 17 | Dependencies |
| Supported the Iceberg connector of the new version | The Iceberg connector of the new version is supported. | 2023-03-17 | Iceberg Data Lake |
| Optimized the job startup message | Before a user starts a job, the system will check whether there are sufficient subnet IPs for the configured parallelism. If no, it will prompt the user. | 2023-03-17 | Job Operations |
| Supported SQL version comparison | Users can compare SQL code and configurations of different versions of a SQL job. | 2023-03-03 | Managing Versions |
| Optimized the Redis connector | The performance of Redis dimensional table query is improved (greatly improved in the scenario of a lot of empty results), and Flink 1.14 SQL jobs can use the Redis connector. | 2023-03-03 | Redis Database |
| Supported the StarRocks connector | The system supports the StarRocks connector based on the open-source starrocks-connector-for-apache-flink v1.2.4. Data can be read from and written to StarRocks via Flink. (Flink 1.13 1.14) | 2023-03-03 | StarRocks Database |
| Optimized MySQL CDC | The system ignores the data changes of non-DDL fields in MySQL CDC. | 2023-03-03 | MySQL CDC Database |
| Supported the MongoDB CDC of the new version | The MongoDB CDC of the new version is supported. | 2023-03-03 | MongoDB CDC Database |
| Optimized cluster purchase | A user other than a super admin can click to purchase a cluster. | 2023-03-03 | - |
| Added display of abnormal jobs and idle cluster CUs | Cluster overview is added to the job management module, which displays the statistics of the clusters associated with this workspace, including total number of clusters, total number of CUs, and number of idle CUs. The total number of jobs with metric alarms and that of jobs with events are added to the job overview section, and support clicking for quick filtering. | 2023-03-03 | Job Overview |
| Optimized the | Users can edit DDL instead of creating a new | 2023-03- | Databases and |

| | | | |
|----------------------------------|--|----|------------------------|
| database/table reference feature | one. They can edit WITH parameters (except connector). A change record button is added to the View DDL page for users to view the DDL change records, which facilitates tracking. | 03 | Tables |
|----------------------------------|--|----|------------------------|

February 2023

| Update | Description | Release Date | Documentation |
|---|---|--------------|---|
| Optimized batch job deletion | The logic of displaying jobs for batch deletion is optimized so that folders containing only jobs that cannot be deleted will no longer be selected by default. | 2023-02-09 | Job Operations |
| Added the SQL Server CDC connector | The SQL Server CDC connector is added. SQL Server CDC can only be used as a source so as to read full data or incremental data from the SQL Server database. (Flink 1.14) | 2023-02-09 | SQL Server CDC Database |
| Supported Doris catalog | The catalog feature of Cloud Data Warehouse for Apache Doris is introduced to avoid manual registration of Doris tables, improving job development efficiency and ensuring correct schema mapping. | 2023-02-09 | - |
| Supported transactions by Doris (data sink) | The system supports end-to-end Exactly-Once for writing to Doris. | 2023-02-09 | - |
| Upgraded the hadoop-shaded package | The flink-hadoop-shaded package has compatibility issues due to a low version (2.7.5) and is upgraded to the latest version 2.8.3. | 2023-02-09 | - |
| Upgraded the KonaJDK version | The KonaJDK version is upgraded to avoid deadlocks. | 2023-02-09 | - |
| Optimized MySQL CDC | The issue that the MySQL CDC string key data is not evenly distributed is fixed so that string primary key chunks are evenly distributed and jobs no longer encounter OOM. MySQL CDC connection merge supports the Kudu data sink. The Kudu connector is refactored with the new Flink Sink API. | 2023-02-09 | MySQL CDC Database |

| | | | |
|---|--|------------|--------------------------------------|
| | The TDSQL GTID issue that caused job failures is fixed. | | |
| Optimized Ingress | The rules of deleting and adding Ingress are optimized. | 2023-02-09 | - |
| Optimized event alarms | The system allows users to configure alarms for "job stopped" events after being redirected to EventBridge. | 2023-02-09 | Events and Diagnosis |
| Optimized job copy creation | After a user creates a copy of a job, the copy can carry a savepoint for fast job migration and restoration. | 2023-02-09 | Job Operations |
| Optimized variable management | A variable name now supports 100 characters. | 2023-02-09 | Variables |
| Supported auto scaling of jobs as scheduled | Users can set the start and stop time, object, and scaling proportion for auto tuning. | 2023-02-09 | - |
| Optimized checkpoint management | The latest five checkpoints are retained for a job by default. Users can configure the number of checkpoints to retain in advanced parameters. | 2023-02-09 | Managing Checkpoints |

January 2023

| Update | Description | Release Date | Documentation |
|---|---|--------------|---|
| Optimized the JDBC (data sink) | Users can configure parameters to determine whether to ignore the Delete operation. | 2023-01-16 | JDBC Connector |
| Optimized the support for Cloud Data Warehouse for PostgreSQL | Cloud Data Warehouse for PostgreSQL (flink-connector-cdw-postgre) v2.0 is supported. | 2023-01-16 | Cloud Data Warehouse for PostgreSQL |
| Optimized the Kafka connector | The incorrect database/table subscription issue is fixed for the Kafka connector. | 2023-01-16 | Kafka Message Queue |
| Optimized the COS connector | The COS connector META-INF Shade error is fixed, and Flink 1.14 SQL jobs can use the COS connector. | 2023-01-16 | FileSystem |
| | | | |

| | | | |
|----------------------------|--|------------|---|
| Fixed TopN operator issues | This issue is fixed: An error may occur due to status mismatch when TopN operators of Flink encounter functions such as now(). | 2023-01-16 | - |
|----------------------------|--|------------|---|

December 2022

| Update | Description | Release Date | Documentation |
|---|--|--------------|---|
| Optimized MySQL CDC monitoring metrics | The following monitoring metrics are added: CurrentFetchEventtimeLag, CurrentEmitEventtimeLag, taskmanager_job_task_backpressuredtimemsperssecond, and taskmanager_job_task_dataskewcoefficient. | 2022-12-16 | Monitoring Metric List |
| Supported automatic table creation in Hive by database sync SQL | Database sync SQL supports automatic creation of missing tables in Hive. | 2022-12-16 | - |
| Supported extraction of metadata fields in database sync | In database sync, metadata fields such as the database name, table name, and timestamp can be extracted. | 2022-12-16 | - |
| Supported more data types for ClickHouse | ClickHouse supports Map, Array, and Row data types. | 2022-12-16 | ClickHouse Data Warehouse |
| Supported Data Lake Compute (DLC) | The system supports writing to native tables managed in DLC. (Flink 1.14) | 2022-12-16 | Data Lake Compute (DLC) |
| Supported TDMQ for Pulsar | The Pulsar SQL connector allows users to query data in Pulsar topics or write data to Pulsar using simple SQL queries or Flink Table API. (Flink 1.13 1.14) | 2022-12-16 | TDMQ for Pulsar |
| Optimized metadata fields | MySQL CDC supports access of metadata fields such as Pos, File, and Gtid. | 2022-12-16 | MySQL CDC Database |

| | | | |
|---|---|------------|------------------------------------|
| supported by MySQL CDC | | | |
| Optimized chunk division by MySQL CDC composite primary key | Users can select a column as the division key in an MySQL CDC composite primary key scenario. | 2022-12-16 | MySQL CDC Database |
| Supported a dynamic key in ZADD by the Redis connector | The Redis connector allows a dynamic field to be used as the key in ZADD mode. | 2022-12-16 | Redis Database |
| Optimized the FileSystem connector | The FileSystem connector supports the compress algorithm. | 2022-12-16 | FileSystem |

September 2022

| Update | Description | Release Date | Documentation |
|--|---|--------------|--|
| Optimized the MySQL CDC performance | The performance for reading incremental data from MySQL CDC is optimized. | 2022-09-16 | MySQL CDC Database |
| Upgraded the COS connector version | The COS connector version is upgraded. | 2022-09-16 | FileSystem (HDFS/COS) |
| Supported custom role permissions | Custom role permissions are supported. | 2022-09-05 | Space Role Permissions |
| Optimized Doris performance | The Doris write performance is optimized. | 2022-09-05 | - |
| Supported Kafka source merge and database sync | The system supports Kafka source merge and database sync (a single topic and multiple schemas). | 2022-09-05 | Kafka Message Queue |
| Supported custom job log levels | Users can customize the log level when configuring running log collection for a job. | 2022-09-05 | Configuring Running Log Collection for a Job |

| | | | |
|--|---|--|--|
| | Users can customize the log level when querying job logs. | | |
|--|---|--|--|