

SingleStore Ingest Operational FAQ

01/08/2025

SingleStore Ingest Operational FAQ	3
How to Restart Ingest	3
Rollback	3
Full Extract	3
Redo Initial Extract	4
Skip Initial Extract	4
Add New Tables to Existing Pipeline	4
Synchronize Schema Changes Using Sync Struct	5
How to Stop Ingest	5
Prerequisites for Oracle Source	5
Create a New Pipeline	6
Latency Calculation	6
Enable Windows Authentication for SQL Server DB	6

SingleStore Ingest Operational FAQ

This document describes some of the most frequent operational user queries.

How to Restart Ingest

Perform the following steps to restart Ingest:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete. Then, stop the Flow Windows service.
2. Restart the Flow Windows service and wait for the Ingest dashboard to appear (refresh the page if needed).
3. Navigate to the **Dashboard** tab, under **Schedule**, enable the scheduler.
4. Ingest picks up the correct processing sequence. Perform a sanity check to ensure that the last sequence number loaded correctly and that the new sequence number started correctly. If needed, perform a **Rollback** to the correct sequence number for processing.

Rollback

Perform the following steps to perform Rollback in Ingest:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete.
2. Navigate to the **Operations** tab, select the gear icon, and then select **Rollback**.
3. In the pop-up window, select the sequence number from which you want to replay the log files. Ensure that the log files are available on the source server.
4. Select **Select**.
5. Navigate to the **Dashboard** tab, under **Schedule**, enable the scheduler.

Full Extract

To perform Full Extract for initial sync of all tables in a pipeline, follow these steps:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete.
2. Navigate to **Dashboard > Tables**:
 - a. Select the tables (one or more) for the Full Extract.
 - b. Ensure the Primary Key (**PKey**) column(s) are applied.
 - c. Enable **Delta Extract (With History)**.
 - d. Select **Apply**.
3. Navigate to **Dashboard > Operations**:
 - a. Select Full Extract.
4. Navigate to the **Logs** tab. Check the logs to ensure that the selected tables have started the Full Extract.

Redo Initial Extract

To perform Redo Initial Extract for existing tables, follow these steps:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete.
2. Navigate to **Dashboard > Tables**:
 - a. Select the table(s) to redo the initial extract.
 - b. Ensure the Primary Key (**PKey**) column(s) are applied.
 - c. Enable **Delta Extract (With History)**.
 - d. Enable **Redo Initial Extract**.
 - e. Select **Apply**.
3. Navigate to **Dashboard > Operations**:
 - a. Select **Sync New Tables**.
4. Navigate to the **Logs** tab. Check the logs to ensure the selected table(s) have started the full extract.

Skip Initial Extract

To perform deltas without initial sync or for XL Ingest, follow these steps:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete.
2. Navigate to **Dashboard > Tables**:
 - a. Select the table(s) to perform deltas without an Initial Sync of the data.
 - b. Ensure the Primary Key (**PKey**) column(s) are applied.
 - c. Enable **Delta Extract (With History)**.
 - d. Enable Skip Initial Extract.
 - e. Select **Apply**.
3. Navigate to **Dashboard > Operations**:
 - a. Select **Full Extract**.
4. Navigate to the **Logs** tab. Check the logs to ensure the selected table(s) skipped the full extract.

Add New Tables to Existing Pipeline

Perform the following steps to add only new tables to existing pipelines:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete.
2. Navigate to **Dashboard > Tables**:
 - a. Select the new tables (one or more) as required.
 - b. Ensure the Primary Key (**PKey**) column(s) are applied.
 - c. Enable **Delta Extract (With History)**.
 - d. Select **Apply**.
3. Navigate to **Dashboard > Operations**:
 - a. Select **Sync New Tables**.

4. Navigate to the **Logs** tab. Check the logs to ensure that the existing tables were ignored and the new table(s) started to sync.

Synchronize Schema Changes Using Sync Struct

To perform Sync Struct for existing tables when the source structure changes:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete.
2. Navigate to **Dashboard > Operations**:
 - a. Select **Sync Struct**.
3. Navigate to the **Logs** tab. Check the logs. Ingest compares the structures and synchronizes any differences.

How to Stop Ingest

Perform the following steps to stop Ingest:

1. Navigate to the **Dashboard** tab. Under **Schedule**, disable the scheduler and wait for the load to complete.
2. Open **Windows Services** and stop the Flow Windows service.
3. If the Flow Windows service stops successfully but the Ingest page is still loading, you may need to terminate the non-responding processes.
4. To terminate the non-responding Ingest process, follow these steps:
 - a. Open the command prompt as **Administrator**.
 - b. Execute the following command:

```
netstat -ano | find "8081"
```

Replace 8081 with the port number on which Ingest is running. For example, if Ingest is running on port 9090, the process ID for Ingest is 4044.

- c. Open **Windows Task Manager**, navigate to the **Details** tab, and find the process with PID 4044.
- d. Right-click on the process and select **End Task**.
- e. This terminates the runaway Ingest Java process.

Prerequisites for Oracle Source

The Flow installation guide provides details on how to enable supplemental logging for source table(s) for Delta replication. Refer to the Flow Installation guide for more information.

Note:

- If database supplemental logging is enabled, only inserts and deletes are captured, and updates are missed.

- If table-level supplemental logging is enabled, only updates are captured, and inserts and deletes are missed.

Therefore, both database and table-level supplemental logging are required.

Create a New Pipeline

Perform the following steps to create an additional ingest pipeline:

1. Unzip the provided file into a folder, such as `C:\singlestore-flow\pipeline\`.
2. Edit the `ingest.xml` file located in the `C:\singlestore-flow\pipeline\` folder.
3. Change the `<id>Bryteflow-Ingest</id>` and `<name>BryteFlow Ingest</name>` tags to reference the new pipeline names.
4. Check that port 8081 is free and not being used by other processes.
5. Open a command prompt as Administrator, navigate to the new pipeline folder, and run `Install-Service.bat`. Then close the command prompt.
6. Open the Windows Services console and start the newly created service.
7. Launch the pipeline at `localhost:8081`.
8. SingleStore recommends changing the port by navigating **Settings > Instance Details > Web Port** to another preferred web port, to avoid using port 8081 for future pipelines.
9. Restart the Windows service and launch the Flow application on the new web port, e.g., `localhost:8082`.
10. Ensure you have a valid license key to get started.

Latency Calculation

Latency is calculated based on the time from when the record is committed to the source until it is available at the destination.

Enable Windows Authentication for SQL Server DB

Steps to enable Windows authentication in Ingest:

1. Stop the Ingest service.
2. Download and copy the `sqljdbc_auth.dll` file to the bin directory of your Java installation. You can download the DLL from [here](#) or obtain it from the Microsoft website.
3. Start the Ingest service and add the following JDBC option in the Source Database connection settings:

```
integratedSecurity=true
```