



FABRIC V6.4.6 RELEASE NOTES

These Release Notes describe the new features in Fabric version 6.4.6 and lists bugs that have been fixed since version 6.4.5.

Certification of this Fabric release is based on:

- Cassandra version 3.11.9 (updated package for Cassandra as a service).
- SQLite version 3.27.2.
- Open JDK version jdk-8u232.
- Confluent Kafka version 5.5.1.
- OrientDB tp3-3.1.2.
- Elastic oss-7.6.0.

MAIN FEATURES AND IMPROVEMENTS

Change Data Capture (CDC)

- Added new optional parameters to the **CDC_REPUBLISH_SCHEMA** command:
 - Tables: Send a list of LU tables for republish to avoid a republish of all the LU tables with CDC fields.
 - Drop_table: This parameter enables a drop and re-creation of the CDC indexes to save a manual drop of the CDC indexes if needed. When the parameter is set to **true**, drop CDC indexes on the LU tables in the list, or all LU tables if the Tables parameter is empty. The default value of this parameter is **false**.

Note: You must run the CDC_REPUBLISH_INSTANCE command on all Fabric LUIs to repopulate the data in the newly created indexes.

- CDC Consumer Creation: No longer requires a partition property when consuming a CDC message from Kafka. The partitions are assigned to the CDC consumers by Kafka.
- Performance Improvements on the Search Loader. Removed the synchronization of the Elastic Search (ED) index data on each data write. The synchronization of the ES index data is performed upon request when running a search command.
- Added the following parameters to the [search engine] section of the config.ini file:
 - BULK_PROCESSOR_MAX_CONCURRENT_WORKERS: maximum number of concurrent threads to process the bulk actions. Default value is 5.
 - BULK_PROCESSOR_MAX_ACTIONS: Bulk size. The maximum number of actions (requests) in one bulk. Default value is 1000.



FABRIC RELEASE NOTES

Common (Reference) Tables

Improved performance when dealing with huge Reference tables (more than 200M records) was achieved by adding a new table property called 'Index Post Sync'. This property determines if the index creation will occur before or after syncing the data. Default value is **false**.

iidFinder

The Solo and Orphan statistics are provided per LU.

RESOLVED ISSUES

- Ticket **22921** – REF_SYNC command is failing on the second reference deployment.
- Fixed the deploy failure when there are multiple CDC jobs that have been stopped before the deploy.
- iidFinder - Close Kafka on exit\exception