FABRIC RELEASE NOTES

FABRIC V8.4.0 RELEASE NOTES

These Release Notes describe the new features in Fabric release V8.4.0 and list bugs that have been fixed since the latest V8.3 release.

Certification of this Fabric release is based on:

3rd Party Name	Included with Fabric	3rd Party Server
Cassandra	java-driver-core-4.17.0.jar	4.1.8
Elasticsearch	8.13.1	8.5.3
Kafka	3.9.1	confluent-8.1.1
Neo4j Enterprise	5.26.19	-
OpenJDK Runtime Environment	openjdk-21.0.10+7-LTS	openjdk-21.0.10+7-LTS
OpenSearch	3.4.0	AWS 1.3.4
PostgreSQL	postgresql-42.7.8.jar	17.5
SQLite	3.46.1	-

MAIN FEATURES AND IMPROVEMENTS

1. Fabric Catalog

The Discovery & Catalog solution has been enhanced with the following management and usability features:

- **Version Reversion:**
  - You can now revert the Catalog from the latest version to any previous state, applied either to a single data platform or the entire Catalog.

  https://support.k2view.com/Academy/articles/39_fabric_catalog/catalog_app/16_revert_catalog_version.html
- **Data Platform Deletion:**
  - Catalog now supports the permanent deletion of a selected data platform, including all its associated metadata across all versions. Note that this action is non-reversible.

  https://support.k2view.com/Academy/articles/39_fabric_catalog/catalog_app/17_delete_catalog.html
- **Granular Artifact Building:**
  - Build artifacts can now be performed for a single data platform or the entire Catalog.
  - A new progress bar in the bottom-right corner displays the progress and the estimated time to completion.

  https://support.k2view.com/Academy/articles/39_fabric_catalog/catalog_app/09_build_artifacts.html

- **Enhanced Version Info:**
  - The Catalog version drop-down now displays the version's origin (Crawler, Revert, or Manual) and a list of impacted data platforms per version.
  - For the existing projects, these details will be populated only for the versions created after the upgrade to V8.4.

  https://support.k2view.com/Academy/articles/39_fabric_catalog/catalog_app/06_catalog_versioning.html
- **Discovery Monitor Enhancements:**
  - New indicators display progress for the crawler, plugins pipeline, and version creation.
  - Steps details are now displayed vertically rather than horizontally, allowing more space for additional information, such as warnings or errors. This section is automatically expanded when there are errors or warnings. It can be collapsed by the user.

  https://support.k2view.com/Academy/articles/39_fabric_catalog/catalog_app/12_discovery_monitor.html
- **Catalog Search Results Enhancements:**
  - The export of the Catalog search results into a CSV file was previously performed by the client application, using the data returned by the search. However, since the search results are limited to 1000, the data export was limited to the same.
  - To bypass the previous 1,000-row client-side limit, the search exports are now performed by the server, enabling the export of all relevant search results.
  - In addition, the Catalog search results screen now includes the **Classification origin** column presenting the plugin name that classified the field.

  https://support.k2view.com/Academy/articles/39_fabric_catalog/catalog_app/08_search_catalog.html
- **Option Set Analyzer Plugin**:
  - When running discovery on JSON Schema files, the plugin now identifies fields with Enum property and saves them in a dedicated Option Set related MTable, for use in masking or synthetic data generation.

  https://support.k2view.com/Academy/articles/39_fabric_catalog/plugins/04_source_data_metrics.html
- **LLM Profiling and LLM Description Plugins**:
  - The default Incremental Mode of these plugins has been changed to 'Keep Existing', meaning the plugin is only invoked for the new fields and for the fields without the 'classification' or 'description' property.

2. Broadway

New Actors

- **PgCopy** actor is a new high-performance actor for PostgreSQL that uses the COPY command for bulk inserts, outperforming the DbLoad actor in batch mode.

Existing Actors

- **RandomNumber**, **RandomDecimal** actors now include a **maxInclusive** input. When set to true, the maximum value is included in the generation range.
    - By default, the maxInclusive input is set to false.
- **RandomDistribution** actor now allows to select an MTable's values for both **weighted** distribution and **constant** distribution.
    - Note that an MTable should be set prior to setting the distribution input param.

3. Web Studio

- **NoSQL Query Builder** enables seamless querying of NoSQL data sources.
    - Users can support NoSQL interface types by creating a Broadway flow named **<interface-type>_query_builder.flow** with predefined input/output parameters.
    - When the flow result is an array of maps, data is presented in table view (as with SQL interfaces). Otherwise, the result is displayed in a Monaco editor panel.
    - Note that several connectors, such as MongoDB, already include the Broadway flow for NoSQL Query Builder support.

  https://support.k2view.com/Academy/articles/11_query_builder/04_query_builder_NoSQL.html
- The Interface Explorer tree now supports the browser native *Copy to clipboard* command, allowing to copy any node name (such as schema, dataset or field).
- **K2exchange:**
    - Searching for extensions is enabled now by default.
    - Performance improvements - load extensions from K2Exchange repository only when Extensions View is chosen.
- **Project Tree** – Improved navigation usability:
    - The schema entry now appears at the top of the LU/Data Product folder.
    - The population files are shown hierarchically under their associated tables at the Tables folder.
    - The previous view can be restored via Studio Settings.

4. Other Improvements and Changes

- The **Parent Rows Grouping** parameter has been introduced in the Fabric properties section of the JDBC interface type definition to enable improved performance tuning.
    - The parameter impacts the WHERE clause generation by the **SourceDbQuery** actor. By default, the parameter is set to **OR** for all interface types except Cassandra (where it is set to **NONE**). Additional options are supported: IN, IN_WITHOUT_TUPLING.
    - When upgrading an existing project to V8.4, if an overridden Cassandra interface type exists in the project – set this parameter manually to **NONE**. For all other overridden interface types, the parameter will be defaulted to **OR**.

  https://support.k2view.com/Academy/articles/07_table_population/15_parent_rows_grouping.html

- **MCP Server support** – Fabric now automatically exposes each LU/data product as an MCP server, including built-in resources and tools, with the ability to easily add custom tools.

  https://support.k2view.com/Academy/articles/46_MCP/01_data-product-mcp-overview.html

- **Security:**
  - **Akeyless** Secrets Management is now supported.
  - **JWT Custom Claims & IID-Based Access Control** - Fabric extends its JWT authentication to support **custom (non-standard) claims** that flow through the session lifecycle. Combined with a new READ_WITH_CLAIM permission, these claims can restrict a session to a single LUI/IID, so that any attempt to read a different instance is denied.

  https://support.k2view.com/Academy/articles/26_fabric_security/06_jwt-custom-claims-and-iid-access-control.html

- **Incremental MDB Export:** Fabric has introduced an incremental mode for MDB exports to optimize data portability and speed.
  - **Selective Export**: By using the INCREMENTAL [version] parameter in the MDB Export command, the system only exports data that has changed since the specified version.
  - This mode significantly accelerates the export process by limiting the scope to modified tables only. To enable this functionality, the configuration parameter TRACK_MDB_CHANGES must be set to true.
- **Object Synchronization Tracking:** The internal **_k2_objects_info** table has been updated to provide better visibility into Logical Unit synchronization.
  - **Change Indication**: Fabric now identifies synchronized LU tables by creating a new entry in this table with the object_name column set to CHANGES.
  - This tracking is only active when TRACK_MDB_CHANGES is enabled (set to true).
- **CDC performance improvement**: The Change Data Capture (CDC) mechanism has been optimized to reduce noise and overhead in downstream systems.
  - **Delta-Only Reporting**: Previously, CDC reported every database write operation. The updated MDB layer now identifies "actual" database changes, ensuring that only true data modifications are captured.
  - This optimization reduces the volume of messages sent, focusing purely on meaningful updates rather than all write attempts.
- **Parallel Entities Transfer to/from S3 and Azure:**
  - Fabric now supports parallel processing for Entity transfers to and from **Amazon S3** and **Microsoft Azure**. The improvement was achieved by splitting large entity BLOBs into chunks for parallel processing, which significantly reduces the large entities' transfer times.
  - To enable this feature, set PARALLEL_DOWNLOAD and/or PARALLEL_UPLOAD to true (both are set to false by default in the config.ini).

- **Web Admin Improvements:**
  - The output of the **test_connection** command now includes the **HOST** column, allowing for faster verification of the target service location.
  - A new **HOST** column has been added to the **Web Admin > Interfaces** tab.
- **OAuth Client credentials authentication flow for Salesforce**:
  - Since the Progress Salesforce JDBC driver does not natively support the client-credentials flow, a wrapper layer was implemented around the driver's OAuth mechanism. This wrapper is responsible for obtaining and managing the access token using the client-credentials flow and injecting it into the driver connection, allowing the driver to operate transparently while complying with Salesforce environments that mandate client-credentials–only authentication. This approach enables integration with Salesforce deployments that require **application-level OAuth authentication without user interaction**.

https://github.com/k2view-academy/K2View-Academy/issues/1459

- The new SET command has been introduced - **SET NODE_AFFINITY**:
  - Allow adding, editing or deleting an affinity from node.id file at run time.
  - The syntax is: "affinity:rec:max" where 'rec' is a recommended number of jobs and 'max' is a maximum number of jobs. For affinity deletion, set 'rec' to 0.
  - For example: SET NODE_AFFINITY "aff1:2:11".
- The **Fabric Reports** application, which uses the third-party Mescius ActiveReportsJS reporting tool, has been upgraded to Mescius V5.2.6.

https://support.k2view.com/Academy/articles/38_reports/01_reports_overview.html

RESOLVED ISSUES

- Ticket #42072 – Web Studio: Enable opening Query Builder from the command pallet.
- Ticket #42779 – Web Studio: The Search view enable directly opening file (no need to reveal it first in the project tree).
- Ticket #45182 — DB interface validation issue / limitation from Prod cluster.
- Ticket #46012 - BuildArtifacts script does not package project-resources folder into the ludbXMLs.zip
- Ticket #46418 – Web Studio: Environment operation is simplified, by enabling set an interface as active/inactive via the interfaces tree rather than entering each interface editor.
- Ticket #46566 – Upsert to DB2 was failing.
- Ticket #47377 – Web Studio: Interface Editor for Kafka now shows properly secret manager options and key buttons.
- Ticket #47417 – CDC message was not sent if IID sync is performed inside the transaction.

- Ticket #47505 – Web Studio: The Query Builder now supports a shortcut to execute the query.
- Ticket #47516 – There was unexpected behavior due to "int8" type in PostgreSQL.
- Ticket #47665 – Random Distribution object was setting NULL when defined as 0.
- Ticket #47755 – Broadway cannot link between 2 fields using the linking tool.
- Ticket #47774 – Integer values converted to Long in MongoDB document when loading to target.
- Ticket #47810 - Web Studio: Environment Globals - if there is override in any LU, the Globals list expanded and appears under each LU, rather than consolidated.
- Ticket #47830 – JsonSchemaToMetadata failed to read a schema.
- Ticket #48024 – Connection to DB2 in pre-production and production was failing with connection exception.
- Ticket #48144 – Some DB2 data types were converted incorrectly in the Catalog.
- Ticket #48201 – Graphit XML invalid characters were handled.
- Resolved 3rd party security vulnerabilities.