



## FABRIC V8.5.2 RELEASE NOTES

These Release Notes describe the new features in Fabric release V8.5.2 and list the bugs that have been fixed since the latest V8.5.1 release.

Certification of this Fabric release is based on:

| 3 <sup>rd</sup> Party Name  | Included with Fabric   | 3 <sup>rd</sup> Party Server |
|-----------------------------|------------------------|------------------------------|
| Cassandra                   | Cassandra-cli-4.1.11   | 4.1.8                        |
| Elasticsearch               | 8.13.1                 | 8.5.3                        |
| Kafka                       | 3.9.2                  | Confluent-8.1.1              |
| Neo4j Enterprise            | 5.26.32                | -                            |
| OpenJDK Runtime Environment | openjdk-21.0.12+8      | openjdk-21.0.12+8            |
| OpenSearch                  | 3.8.0                  | 3.1.0                        |
| PostgreSQL                  | postgresql-42.7.13.jar | 18.4                         |
| SQLite                      | 3.46.1                 | -                            |

## MAIN FEATURES AND IMPROVEMENTS

### 1. Studio

- AI:
  - **K2assistant chat modes:** Each chat session now has a mode selector: Ask, Plan, Act (default)
  - **Interrupted sessions resume:** After a page reload, a disconnect, laptop sleep or a backend restart, Studio notifies user and points him to interrupted chats, suggesting him to resume them.
  - **Keep chat working on long sessions:** Let users, with a click, keep long sessions working, by keeping the screen awake while those chats are running.
  - **New memory capability:** K2assistant can keep a wiki-style knowledge base in the workspace, turning what it learns in a session into linked notes.
  - **K2assistant and K2worker are now built into Studio:** They no longer load from the Studio AI core artifacts extension. Those agents which are left over from older setups (Studio 8.5.1 and earlier) are ignored.
  - **Confirmation before sensitive AI tools run:** The assistant now asks before it runs some tools that might contain sensitive data, like tables data and schema info.



# FABRIC RELEASE NOTES

- **AI Configuration reorganization and facelift:** AI Configuration contains now all AI settings in a single view with a category tree list.
- **Search within a chat session** is now available. (CR Ticket #51406)
- Agents are now warned when a file they read has changed since, so they don't overwrite user's recent edits.
- Anthropic requests no longer fail on invalid "thinking" blocks in the chat history.
- The model list is now read from each provider dynamically and it is no more a fixed list setting.
- **AI walkthrough:** Introduce new onboarding steps for using AI within Studio as part of the welcome tab (which is now shown on every reload once there are no opened tabs) and it now contains a new AI walkthrough component.
- AI model aliases updated to the latest available models.
- **Deploy checks for Java compilation errors.** Before any deploy done by user, Studio checks for Java errors in the Data Product or Shared Objects code, and if it finds any, it notifies user about and suggest if to proceed with deploy or skip. (can be turned on and off at Settings)
- **Exchange:**
  - Uninstall an extension now does not delete imported files which are shared/used by another installed extension.
  - Pre-release extension versions are now hidden from the Exchange list and search results (unless the "show pre-release extensions" setting is on). Such versions are now also marked in the version picker. This replaces the "Preview" sign, which remains as a label, on extension level, not used for filtering.
- The project tree no longer jumps back to the selected item while you scroll.

## 2. Catalog

- The discovery /catalog/start-crawler-job API can now run in an **ad-hoc mode** for the specified data platform and schema. In ad-hoc mode:
  - When a schema name is provided in the API input, the discovery is executed only for this schema, while all other schemas of the same data platform are excluded from the run.
  - This ad-hoc rule is combined with the existing rules (in pluginsOverride.discovery) however it is not saved as a new rule.
  - Ad-hoc mode is not applicable when discovery is triggered by the Catalog App.
- The LLM Domain plugin can now work in incremental mode.

## 3. Miscellaneous

- **JDBC Prepared Statement Handling** - updated the default JDBC prepared statement configuration to improve handling of `NULL` values, including `BLOB` fields, by using explicit JDBC type information when setting `NULL` parameters.



# FABRIC RELEASE NOTES

- Certified Fabric with Cassandra 5.0.9.
- **STARTJOB** command: added a new optional Boolean parameter, `RUN_AS_USER`. When set to true, the job records the user who started it in the current session. When omitted or set to false, the job's user field remains null.
- **CDC Search: Configurable Elasticsearch Index Settings**
  - Fabric now lets you customize the settings of the Elasticsearch indexes that the CDC Search consumer creates, such as shard and replica counts, using external JSON settings files.
    - **Root settings:** A JSON file placed directly under `workspace/config/cdc-tags/` holds the root settings, which apply to all environments.
    - **Environment settings:** You can add a file with the same name under `workspace/config/cdc-tags/<environment>/` to customize settings for that environment. In Kubernetes, you can supply these files as ConfigMaps.
    - **Merge behavior:** The two files are merged:
      - Attributes that appear in only one of the files are combined.
      - When the same attribute appears in both files, the environment-folder value wins.
      - The combined result takes precedence over system defaults.
    - **Per-index configuration:** You can define a separate settings file for each index.
    - **Applying settings:** The merged settings are included in the index create/update messages that `CDC_REPUBLISH_SCHEMA` sends.
    - **Graceful handling:** If a settings file is missing, Fabric ignores it and uses the remaining configuration or the system defaults.
  - **Note:** The operator maintains the content of each environment's file and is responsible for making sure a change is compatible with the existing index. Some settings, such as replica count, can be updated on an existing index. Others, such as shard count, can't be changed after the index is created. To change those, reshard the index first, then update the settings file.

## RESOLVED ISSUES

- Fixed a memory leak in the `iidFinder` job.
- Fixed the `SourceDbQuery` cache key so generated queries are now cached and reused instead of being re-parsed on every call.
- Fixed an issue where `error.graphit` returned HTTP status code 200 even when the underlying logic failed.
- Fixed an issue in the Catalog which occurred when running the Bulk Dimension Loader plugin on nested / complex classes.



# FABRIC RELEASE NOTES

- Fixed an issue in the Studio Interface Explorer which prevented displaying the fields upon expanding a table, when a field name is in Hebrew.
- Fixed a regression issue in the /interface\_schema/table/ API which returned wrong PKs.
- Fixed an issue when setting `global affinity_rules= ''` did not stop an Affinity job when the DC name contained a `-`; the job is now correctly stopped when the affinity rules are empty.
- Fixed a `NullPointerException` thrown by a custom Catalog plugin which extended the Relation Plugin, when the custom plugin didn't have any input parameters and a threshold definition.
- Ticket #50939 - improved synchronization when retrieving topics from the cache to prevent race conditions in which multiple threads could acquire the same topic concurrently.
- Ticket #51165 – for Cassandra lightweight transactions, a `CASWriteUnknownException` will trigger up to three retry attempts before the operation fails.
- Ticket #51275 – Minor GUI defects in Catalog Sequence screen.
- Ticket #51289 – prevent users from assigning a different classification to a field in Catalog if it is already classified as a sequence.
- Ticket #51319 - SELECT query returned different values in Fabric Terminal vs. Admin GUI for an Integer column containing decimal values.
- Ticket #51381 - Studio AI execution on space with timeout on browser (see above 2 new capabilities of resuming sessions and keep long chat working).
- Ticket #51439 - removed the Environment folder and its environment files from the LU deployment process, as they were unnecessarily included as Shared Objects and are not required for LU functionality.
- Ticket #51548 – Studio AI - Chat window scrolls up while typing.
- Ticket #51595 – Discovery job creates a new version without any changes after manual overrides are applied on the Catalog.
- Ticket #51659 – resolve rounding issues in Broadway when handling large numbers.
- Ticket #51743 – Catalog doesn't save new rules after upgrade.
- Ticket #51795 – Unable to add a property to field inside the complex structure.
- Ticket #51859 - Mongo connector issue with selection of template for Data product.
- Ticket #51642 – GCS interface: added a new Credentials Content (JSON) field alongside the existing credentials file field, with Secret Manager support. If both are set, the JSON content is used first, and the file is used as a fallback.
- The following 3rd Party tools were upgraded:
  - Neo4J – 5.26.32
  - Studio Java plugin – 1.56
  - Theia - 1.76



# FABRIC RELEASE NOTES

- bc-fips – 1.0.2.7
- Jline-builtins - 4.3.1b
- netty-codec-dns – 4.1.136
- jackson-databind-2.18.10 and 3.1.6
- mssql-jdbc 13.6.0
- zstd-jni 1.5.7\_21
- lz4-java 1.11.4