



TDM V9.4.0 RELEASE NOTES

These Release Notes describe the new features in TDM release V9.4.0 and list bugs that have been fixed since the V9.3.x release.

- Certification of this TDM release is based on:
 - Fabric 8.3
 - PostgreSQL 17
 - MongoDB extension 1.0.14
 - Couchbase extension 1.3.0

MAIN FEATURES AND IMPROVEMENTS

1. Native Support for NoSQL Data Sources

Support for Complex Documents

- TDM V9.4.0, which is based on Fabric V8.3.0, now offers native E2E support for NoSQL Document Storage (such as MongoDB or CouchBase). This support includes Catalog discovery, Catalog masking and sequence handling, parsing a complex document into an LU, and load and delete capabilities for these documents:

Discovery:

- A Discovery job can run on interfaces such as MongoDB or CouchBase, once the corresponding K2exchange connector has been installed in the project. The Catalog is then created based on the discovered document hierarchy. (This feature was already supported prior to V8.3).

Logical Unit implementation:

- The Web Studio's Interface explorer can now display complex document structures, such as nested hierarchy levels and arrays of primitives.
- Logical Units can now be created based on the Document's metadata retrieved from the Catalog, after Discovery process has run on it. Nested hierarchy levels are then created as LU tables, each with a referential link to its respective parent level.
- Root table population is the only action that reads data from the source (the Document). The remaining populations receive the related data, extracted by the root table population, and use it to populate the child LU tables.

For more information about these steps, see [Fabric V8.3.0 Release Notes](#).



TDM RELEASE NOTES

Load, rule-based generation, and delete flows:

- TDM V9.4.0 has been enhanced to support LUs that are generated based on NoSQL documents and enable loading masked documents into the target environment. The *TDMLUInitBasedOnFabric* flow now supports various interface types, including NoSQL connectors, and adds the TDM setup to LUs that are created based on complex documents. The following enhancements have been added for complex documents:
 - Adding a TDM setup to the LUs.
 - Updating the generated load flows for complex documents — a single load flow now assembles the complex document structure based on the LU tables and loads the entire document.
 - Updating the generated delete flows for complex documents — a single delete flow now deletes the document from the target environment.

Supporting NoSQL Data Sources

- Previous TDM versions automated the TDM implementation for JDBC data sources. TDM has now been enhanced to automate implementation for NoSQL data sources as well. The relevant templates are included in the data source connector's extension. The *TDMLUInitBasedOnFabric* flow now supports adding the TDM setup to LUs that are based on NoSQL data sources.

2. Allow Testers to Run Tasks on All Matching Entities (tickets #40559, #42349 and #41950)

- This feature adds the option to grant the user permission to process unlimited number of entities in a task.
- Users with this permission can leave the *Max number of entities* field empty when creating tasks based on business parameters or custom logic selection methods. The task will then run on all entities that match the selected criteria.
- The *Predefined entity list* selection method is now available for tester users with a permission for unlimited number entities.

3. TDM Self-Service — Table-level Tasks Enhancements (ticket #43908)

- Enable the user to select all tables in a given DB schema instead of selecting them one by one.
- By default, the available interfaces are displayed without their related table names. The user can select an interface to view its schemas and select a schema to view its tables. This way, only the relevant information is displayed in the window.



4. Task Creation — Adding a Description on Business Parameters (ticket #35955)

- The implementor can now add a short description to the predefined business parameters.
- The description will be displayed in the task when hovering over a parameter, which would help the user to select business parameters for entity sub-setting in the task.

5. Task Execution — Adding an Execution Note to the List of Successful and Failed Entities (tickets #42601 and #42211)

- Adding an execution note at the entity level. This enables the project implementor to populate execution notes on processed entities.
- The task executor can view the execution notes in the task execution window and receive additional information about executed entities.

6. Task Execution and Reserved Entities Windows — Enhanced Entity Display (ticket #42599)

- Improved the display of processed entities in both the Task Execution Detailed Statistics and Reserved Entities self-service windows. Full source and target entity IDs are now fully visible regardless of their length, and users can copy the complete IDs as needed.

7. Persistent View on Tasks Window (ticket #42600)

- Preserving the user's filters on the task list when navigating to other self-service windows.

8. Concatenating the Task Name to the Batch ID Upon Task Execution

- The task name (used as a prefix; up to 18 characters allowed) is concatenated to the batch ID to simplify tracking of batch processes initiated by task executions. This prefix is added using the new parameter — *BATCH_ID_PREFIX* — of the Fabric 8.3 *BATCH* command.

9. Rule-based Data Generation Parameters Display (ticket #42734)

- All rule-based data generation parameters are now visible when opening a task, regardless of whether they use default values. Previously, only parameters modified from their defaults were displayed, making it difficult to view the complete parameter selection.



10. Additional Fabric V8.3 Related Features

- The Discovery & Catalog solution has been enhanced with several powerful features, such as:
- The ability to create a bulk of Catalog nodes and edit this bulk (e.g., add new property to all bulk entries simultaneously). The bulk edit allows users to efficiently update multiple Catalog entities, by adding or editing various properties simultaneously. This feature simplifies the Catalog sequence setup and enables assigning a Sequence Name property to multiple table fields.

For more information about this feature, see [Fabric V8.3.0 Release Notes](#).

11. SDG Based on AI — Improvements

- Added new execution parameters to improve performance.
- Added support of tables without PKs. Note that the schema must still include physical FK relations between the parent-child tables.
- Added a new configuration table — **AIConfigParams** — to hold the infrastructure configuration parameters for the AI processes.
- **Evaluation Enhancements:**
 - TDM self-service application — a new checkbox, Evaluate generated data, has been added for AI-based data generation tasks. When checked, the task execution generates an evaluation report on the generated data.
 - Improved performance of synthetic evaluation tasks, significantly reducing execution time.
 - Added metrics to the evaluation report.
 - Improved handling of unsupported, empty, and null-valued columns.
 - Corrected issues relating to web chart display, metric propagation, and spider chart score calculations.
 - Dynamic report visuals:
 - Spider charts are displayed only when metrics are available; otherwise, bar plots are shown.
 - Each category's pie chart is now paired with its corresponding metrics plot.
- **LRI (Logical Relations Identification) Enhancements:**
 - Introduced hyperparameters to control or reduce excessive relation detection.
 - Optimization of logic when working with small datasets and composite keys.
 - Improved logging and recognition of logical relations.
- **Metadata Improvements:**



TDM RELEASE NOTES

- Added support for constant columns per PK and for the UUID data type.
- Additional metadata included to support debugging transformations applied to GPT model inputs.
- **Special Fields Handling:**
 - Improved handling of columns with millions of unique values.
 - Corrected Special Field initialization and post-processing logic.
- Refined logic for training, joining and sampling.
- Fixed multiple edge cases that cause crashes during patching, evaluation and generation.
- Addressed issues with NULL values and unsupported data types.
- Corrected handling for Boolean and Timestamp data types.

12. Miscellaneous

- TDM self-service application —
 - Enabling the user to populate a value that contains a comma for a task variable or an environment variable (tickets #41963 and #44059).
 - Adding alphabetical sorting of Logical Units (LUs) in the drop-down list when adding an LU to a BE TDM GUI — sort the LUs alphabetically, as it makes adding LUs to a Business Entity (BE) or a System easier (ticket #42230).
 - Overlapping elements in Task window — preventing the task's widget from overlapping with the 'Advanced settings' when using a small window (ticket #42256).
 - Task Execution — Detailed Statistics window — updating the status from Copied to Successful.
 - Preserving tasks when removing pre or post execution processes from the Business Entity (BE) — the removal of pre or post execution processes from a BE now automatically removes them from the related tasks and the tasks remain active (ticket #43437).
 - Task execution summary window — display an error message if no entity is found for the task execution. Previously, this error was included only in the execution report.
- Enhancing the TDM templates to enable adding tables with alpha-numeric column names to an LU (ticket #43509).
- Parameters coupling mode — performance improvement (ticket #42605).
- Enhancing the TDM upgrade process to simplify the upgrade and avoid overriding the setup of TDM Globals:
- Split the TDM shared Globals into two files —



TDM RELEASE NOTES

- Core Globals — this file includes internal TDM Globals that should not be overridden by the implementation; the file is overridden by the TDM upgrade.
 - Shared Globals that can be edited by implementation (for example, TDMDB_SCHEMA Global). The TDM upgrade merges the new Globals (if exist) into this file but does not override the existing values.
-
- Automating removal of TDM relation tables from LUs (ticket #44230).
 - The TDM DB creation now retrieves the user DB from the TDM interface, if it exists (ticket #41403).
 - Enhancing the TDM deployment flow to run a conditional creation of the TDM DB (create the TDM DB only if the CREATE_TDMDB Global is set to true). This is required if the user must update the default TDM DB schema name, as set in the TDMDB_SCHEMA Global, before creating the TDM DB (ticket #42823).

RESOLVED ISSUES

- Task execution report - statistics report —
- Fixing the statistics report on extract tasks (ticket #44039).
- Fixing the error that occurs when an LU table includes multiple populations (ticket #42760).
- Empty tables are included in the statistics report and shown with a count of zero (i.e., count=0) (ticket #43620).
- Fixing the report to exclude failed entities (ticket #43022).
- Fixing the task execution with multiple post execution processes (ticket #43327).
Fixing the task execution report — the length of error message fields has been limited in order to prevent Excel corruption when downloading the report from the TDM self-service application (ticket #43879).
- Fixing a mismatch in execution status results, where the Execution Summary window and the TDM execution report showed different outcomes (ticket #43212).
- TDM templates —
- Fixing the TTL parameter for the Masking Actors (ticket #43764).
- Fixing the *TDMLUInitBasedOnFabric* flow, which was generating an empty column in the target LU tables (ticket #42133).
- Execution Summary window —
 - Fixed the URL used when opening the Batch Monitor window from the Execution Summary window (ticket #42739 and #42988).



TDM RELEASE NOTES

- Disabled the List of Entities icon while the task is running in order to prevent an incorrect syncing of the task execution into the TDM LU (ticket #43228).
- Task window —
 - The 'Do not delete' option for a retention period is no longer displayed if 'Create data snapshot' is checked and the tester is not allowed to save data for an unlimited period for a newly created data snapshot.
 - Fixing the issue where the task window failed to correctly calculate the number of matching entities based on the business parameters set by the user, as numeric parameters were not treated as numbers. This issue happened when using the regular parameters handling mode (ticket #44098).
 - The Load and Delete checkboxes are now pre-set as checked and disabled for editing when the task's *Policy for fetching data* is set to *Selected snapshot* (ticket #43235).
 - Fixing the warning validation on reserved entities.
 - Fixing the display of valid values for business parameters to prevent duplicate values if the parameter's Logical Unit (LU) is included in multiple Business Entities (BEs).
 - Refreshing the number of matching entities when a selected business parameter is deleted from the task.
- Fixing the error that occurs when manually running a load flow in a debug mode.
- Table-level task —
 - If there is no schema (e.g., CSV file) or the schema is dynamic, TDM now retrieves the information for the task's schema from the Catalog (tickets #43843 and #43984).
 - Table Level supports all field types supported by SQLite, such as DECIMAL, DATE (ticket #43524).
 - Fixing the setting of the source environment for all tables to maintain data consistency of masked PII fields when the data consistency level is defined per environment in the Catalog.
 - Fixing the error that occurs when running a table-level task with a dynamic schema (ticket #42663).
 - Fixing the table filter applied during the execution of extract-and-load tasks to filter only source values.
- Fixing the overriding execution parameters when executing a task via API —
 - Fixing the parsing error that occurs when overriding the number of entities (ticket #44062)



TDM RELEASE NOTES

- Validating JSON for overridden execution parameters when running the start task from API (ticket #44212)
- Supporting parameter names that include spaces for custom logic flows (ticket #43033).
- Fixing the error that occurs when invoking the *wsTaskMonitor* API (ticket #43219).
- Fixing the error that occurs during synchronization of task execution of entity clone into TDM LU. This happens when implementing the sequence handling based on the Catalog (ticket #43006).
- TDM Self-Service — catching a timeout error and giving a comprehensive error message to the user (tickets #42822 and #42110).
- Fixing the error handling mechanism applied during the execution of rule-based data generation tasks.
- 3rd party high and medium security vulnerabilities resolved.