



TDM V9.3 RELEASE NOTES

These Release Notes describe the new features in TDM release V9.3 and list bugs that have been fixed since the V9.2.x release.

- Certification of this TDM release is based on:
 - Fabric 8.2.1 HF1
 - PostgreSQL 17
 - BigQuery extension 1.3.2

MAIN FEATURES AND IMPROVEMENTS

1. Supporting Table-Level Task with BigQuery Connector (ticket #41125)

- The BigQuery connector provides integration with Google's BigQuery platform. Both the BigQuery connector and the TDM library have been enhanced to support the TDM table-level tasks on Google's BigQuery platform.
- Note that data retention is not supported when the source tables contain ARRAY, STRUCT, or RANGE complex fields. The *'Retention Period'* field must be set to *'Do not retain'* when creating a table-level task for such tables.

2. Vertical Execution – Performance Improvement

- Enabling parallel processing of child entities for each root entity. This improvement is based on the new Fabric 8.2 Actor - InnerFlowSession - that opens a new Fabric session and enables running an attached inner flow asynchronously, either locally (default) or on a random node.

3. AI-based Entities Generation – Improvements

- **Special Field Model Enhancements:** Introduction of minGRU (minimal Gated Recurrent Unit) architecture; optimization of training batch size, number of layers and trainable parameters; implementation of early stopping.
- **Vocabulary and Model Efficiency:** Reduction of GPT's vocabulary size by mapping special fields as Booleans; decrease of embedding size, optimization of hidden state size; reduction of the number of attention heads and layers.
- **Dataset Optimization:** Application of various techniques such as sampling, mapping, as well as data type changes for the purpose of optimizing input datasets based on the used model.
- **Logical Relations Optimizations:** Improvement of dataset sampling sizes for identifying and training logical relations; significant increase of batch size for generating logical relations.



TDM RELEASE NOTES

- Introduction of a **new post-execution process for data generation tasks** - **EvaluationDataSubset** - that assesses synthetic data quality by generating two key outputs:
 1. An **Evaluation JSON File**, which consolidates scores based on metrics such as distribution consistency, data validity and novelty.
 2. An **Evaluation HTML Report**, which presents results with interactive graphs, plots and tables, enabling easier interpretation of generated data quality.

4. TDM Environment Management Improvements

- Ticket #40645 – enabling either admin or environment owner user to disable system(s) on a TDM environment. Disabling a task execution when a task needs to run on disabled systems.
- Limiting some users to use only pre-extracted entities from an environment instead of extracting entities themselves. If the user does not have a permission set on the source environment, the user could still get pre-extracted entities from this environment and load them to their environment.

5. TDM Execution – Performance Improvements

- Merging TDM performance improvements from V9.1.5 and V9.1.6 into V9.3:
 - Adding the following changes to the TDM LU in order to reduce its size and optimize performance for large-scale tasks that involve hierarchical Business Entities structures:
 - Eliminating data duplication in the TDM LU.
 - Adding an optional mode that excludes the entity hierarchy in TDM LU and prevents duplicate record occurrences when child entities have multiple parent relationships.
 - Optimizing the population and update of task_execution_entities and task_execution_summary TDM DB tables.

6. Adding a Masking-Only Mode to the TDM Portal

- A *Masking-only* mode has been added to the TDM Portal. When no Business Entities (BEs) are implemented, the TDM Portal now automatically operates in *Masking-only* mode, which supports table-level tasks exclusively.
- To enable the full functionality with BE-based tasks:
 - Deploy at least one Logical Unit (LU) to activate the Business Entities and Systems tabs.
 - Define at least one Business Entity (BE) to enable the creation of BE-based tasks.



RESOLVED ISSUES

- Ticket #40445 – fixing of table-level task execution: Retrieve the tables’ FKs from the JDBC driver instead of from the Catalog in order to set the tables’ execution order. This fix is required due to the following:
 - Non-dev environments do not have the Neo4j schema with the tables’ relations.
 - The Catalog’s identified relations include physical and logical FKs. The tables’ execution order should be set based on the physical FKs.
- Ticket #41585 – fixing of the Reserved Entities window – clicking the Esc button used to behave as a confirmation and execute the escaped activity. Now resolved.
- Ticket #41586 – fixing of the TDM template for the LU population for the main source LU table – the CatalogMasking Actor settings is fixed.
- Ticket #41581 – fixing of the TDM template for the LoadAllTables flow –a condition added for checking the load status for each flow’s Stage.
- Ticket #41565 – fixing of the TDMLunit flow that adds the TDM setup to the LU – set the target environment as the active environment before creating the TAR tables.
- Ticket #41583 – supporting the use of a given flow as pre and post execution process in a task.
- Ticket #40813 – removal of the TDM Demo Globals from the TDM Library.
Ticket #42115 – fixing of the number of matching entities’ calculations when creating a task based on business parameters: adding a filter out of reserved entities of task executions that have been cleaned.
- Ticket #42116 – fixing of the Reserved Entities window to include also valid reserved entities of task executions that have been cleaned.
- Ticket #41715 – optimization of the API that gets all the tasks for the Tasks window in order to have better performance when having a large number of Inactive tasks. The API now gets by default only the Active tasks. If the user changes the filter on the window, the API is invoked again to extract the tasks based on the updated filter.
- Ticket #41961 – fixing of the checkbox checking to mask the task’s PII fields when executing a task based on the task’s environment’s definition.
- Ticket #41042 – the permission group mapping window –
 - Adding a warning message to the user when changing his own role from Admin to Tester/Owner.
 - Do not allow a change of the permission group type from Admin to tester/owner if there is no additional admin role.
- Ticket #39749 – task creation – the Load option is disabled if the user has a permission to only reserve the entities on the target environment.



TDM RELEASE NOTES

- Ticket #42093 – fixing of the TDM upgrade process for TDM versions older than 9.0.
- Fixing of the task execution based on a predefined list and MySQL DB type.
- Table-level – fixing of the list of available selected snapshots (versions) for the tables to check the interface and schema in addition to the table name. This fix is needed when different interfaces or schemas have identical table names.
- Fixing of the execution of rule-based generation when setting a custom separator for the entity IDs (LU instances).
- Parameters coupling parameters – fixing of the API that calculates the number of matching entities when using multiple parameters with *OR* operator.
- Task execution report – fixing of the display of the custom logic's input parameters.

- TDM Portal fixes –
 - Population fixing of the pre and post execution processes' input parameters in the task.
 - Fixing the way that a task is saved in cases where the user edits and then clears the Target component.
 - Task creation - the number of entities is now automatically set to one in the Subset component for the entity clone task.
 - Removal of the Batch Monitor icon when executing a table-level task.
 - Do not display the 'Do not Delete' option in a data snapshot creation task for testers if they are not permitted to store the data snapshots for an unlimited period.
 - Rule-based data generation task –fixing the display of available executions within a date range when loading pre-generated entities.
 - Environment – enabling the creation of a reserve-only permission set when the environment type is Both, i.e., the environment can be used as both source and target environments.
 - Fixing the number of entities in the task execution summary report and window for entity clone tasks.
 - Fixing an issue in rule-based generation tasks where returning to the Source component, after configuring other components, would accidentally reset previously entered Source parameters. Users can now freely navigate between components without losing their Source configuration data.
 - Load data version task – fixing of a typo in the execution note on displayed data versions.
 - Enabling users to copy displayed error messages for failed entities.



TDM RELEASE NOTES