



trino

# The Open Source journey of the Trino Delta Lake connector

**Boston Trino Fest** June 2024.




**Marius Grama**



**Delta Lake**

# About me

- Marius Grama



Starburst  
ANALYTICS ANYWHERE



- 



[findinpath](https://github.com/findinpath)



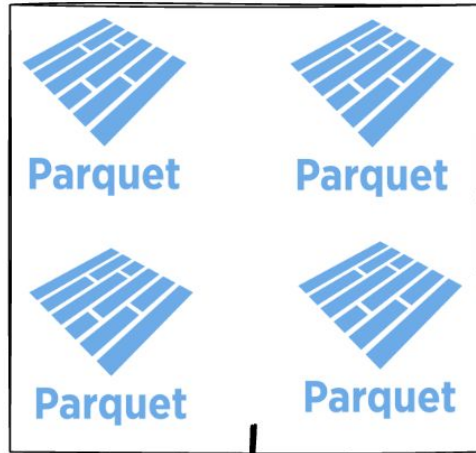
trino

# Delta Lake Table Format


# Contents of a Delta table



DELTA LAKE



Data stored  
in Parquet files



Transaction log  
with metadata

# Contents of a Delta table

```
mytable/date=2020-01-01/1b8a32d2ad.parquet
├── /a2dc5244f7.parquet
├── /date=2020-01-02/f52312dfae.parquet
│   └── /ba68f6bd4f.parquet
└── /_delta_log/000001.json
    ├── /000002.json
    ├── /000003.json
    ├── /000003.parquet
    ├── /000004.json
    ├── /000005.json
    └── /_last_checkpoint
```

Data objects  
(partitioned  
by date field)

Log records  
& checkpoints

Contains {version: "000003"}

Transaction's operations, e.g.,  
add date=2020-01-01/a2dc5244f7f7.parquet  
add date=2020-01-02/ba68f6bd4f1e.parquet

Combines log  
records 1 to 3







trino

Open sourcing Trino Delta Lake connector



# Open sourcing Trino connector

## - Context

-  **databricks** creates [#10300](#) *Venki Korukanti* December 2021
-  **Starburst** ANALYTICS ANYWHERE responds by open sourcing their connector [#10897](#) *Jakob Komissar* February 2022
- Agreement between parties leads to [#10897](#) being contributed to 
  -  **Denny Lee** is supportive and offers cloud testing resources to Trino



trino

# Features Journey





# Beginnings in the open source

- Support for new DDL commands
  - ALTER TABLE ADD COLUMN [#12371](#) Yuya Ebihara May 2022
  - COMMENT [#12452](#) [#12971](#) Yuya Ebihara May 2022
- Support reading tables having `delta.columnMapping.mode=name` [#12675](#) Yuya Ebihara June 2022
  - relevant for schema evolution scenarios
- Increasing interest in the community about new Delta Lake features

# MERGE

- MERGE support [#7933](#) *David Striker & David Philips* August 2022

- dbt-trino  **dbt**  adopts it swiftly [#104](#) through *Przemysław Denkwicz*
  - *incremental* materialization
- Interest on both **Trino** & **dbt-trino** grows significantly

# Second half of 2022

- Support ALTER TABLE RENAME TO [#13707](#) *Konrad Dzedzic* September 2022
- Support reading tables having `delta.columnMapping.mode=id` [#13678](#) *Yuya Ebihara* September 2022
- Parquet performance improvements [#14247](#) *Raunaq Morarka* September 2022, ...
- Support register\_table [#13568](#) *Vikash Kumar* November 2022
- Support for Trino views [#11763](#) *Michiel de Smet* December 2022



# S3 related improvements

- S3 directory/prefix deletion with batch requests [#13974](#) *Marius Grama* December 2022
  - Identified when noticing that deleting managed Databricks table leaves empty objects behind
  - Massive help from *James Petty* to create efficient S3 specific algorithm to batch the deletions
  - First hints from *David Philips* about decoupling Trino from Hadoop file system
    - Implementation of [#15921](#) started on February 2023

# First half of 2023

- Support writes in tables with CHECK constraints
  - [#15396](#) Yuya Ebihara February 2023
  - [#17120](#) Stawomir Pajak April 2023
  - [#17195](#) Yuya Ebihara July 2023
- Pushdown reading of nested columns [#17085](#) Vikash Kumar May 2023
- DML support on tables with name and id column mappings
  - [#16183](#) Yuya Ebihara April 2023
  - [#16600](#) Marius Grama May 2023
  - Add support for writer version 5
- Change Data Feed support via table\_changes [#16205](#) Konrad Dziedzic May 2023
  - Support for polymorphic table functions added through Kasia Findeisen & Martin Traverso July 2022

# Second half of 2023

- Schema evolution support
  - Support for dropping columns [#15792](#) Yuya Ebihara July 2023
  - Support for renaming columns [#15821](#) Yuya Ebihara July 2023
- Support for reading Deletion Vectors [#17477](#) Yuya Ebihara September 2023
  - Merge on Read
- Add support for writer version 7 [#18423](#) Yuya Ebihara September 2023

# Checkpoint filtering `delta.checkpoint-filtering.enabled`

- *Jonas Kyling* has been initially pushing improving the efficiency for reads from checkpoint files [#17408](#) July 2023
  - Initial inspiration for checkpoint filtering functionality
- Improve performance on reading large checkpoint files
  - Driven by *Piotr Findeisen & Raunaq Morarka*
  - [#19848](#) prune unused stats columns from the checkpoint file *Marius Grama* November 2023
  - [#19588](#) partition pruning when reading the checkpoint file *Yuya Ebihara* November 2023

# First half of 2024

- Support writing concurrently to disjoint partitions [#18506](#) *Marius Grama* February 2024
- Enable checkpoint filtering [#20901](#) *Marius Grama* March 2024
- Support reading V2 checkpoint files [#19507](#) *Yuya Ebihara* March 2024
- Support CREATE OR REPLACE support [#19991](#) *Praveen Krishna & Jonas Kyling* April 2024
- Support Time Travel by version [#21052](#) *Yuya Ebihara* April 2024



# Local file system caching metadata & data files

- *Jonas Kyling & Florent Delanoy* address concerns about their AWS S3 bill by reducing the metadata reads for Delta Lake tables
  - Add the ability cache files on the local SSD [#18719](#) February 2024
  - *Raunaq Morarka* guiding force

# Integration with Databricks Unity metastore

- Support for using HTTP as transport for the Thrift metastore client [#20371](#)

*Vihang Karajgaonkar* February 2024

- Trino can connect with Databricks Unity to read data from *external* tables

# CREATE OR REPLACE

- CREATE OR REPLACE adoption driven & implemented by *Michiel de Smet* [#13681](#) October 2023
- CREATE OR REPLACE support [#19991](#) *Praveen Krishna & Jonas Kyling* April 2024
  - Add logic to writes checkpoint if necessary to [#21609](#) *Belen Pruvost* May 2024
    - reduce eventually the query analysis time

# Time Travel

- Highly requested features since the inception of the connector
  - Workaround: create SHALLOW CLONE of a past version of the table
    - [#17038](#) *Vinay Bharadwaj* April 2024 adds support for absolute paths
- Time Travel by version added through [#21052](#) *Yuya Ebihara* April 2024

# Wizards of the Delta Lake connector

- Alex Jo
  - Wrote the initial version of the connector
  - Reviewed most of the important functionality added on the connector
- Piotr Findeisen
  - Shaped & reviewed the important functionality of the connector
  - Constant outstanding improvements in the code quality and accuracy of the connector
  - OSS Champion
- Yuya Ebihara
  - Explosive amount of proactive contributions in making Trino a real alternative to Spark engine for dealing the Delta Lake workloads
  - OSS Champion
- **Trino community**



trino

Next?


# New potential features

- Deletion vectors write support [#22102](#)
- Store table metadata on the metastore [#21463](#)
  - Accelerate information\_schema.columns & system.metadata.table\_comments queries
- Better integration with the Unity catalog
- Liquid Clustering



trino

# Questions?





Thank you!

