

SharePoint Lists specific

Introduction


Microsoft Fabric provides native support for Delta Lake tables within Lakehouses, including time travel and versioning capabilities. However, when working with **mirrored tables (CDC via Mirroring)**, these capabilities are not directly exposed in the Fabric UI as they are in standard Lakehouse tables.

This document summarizes how to access:

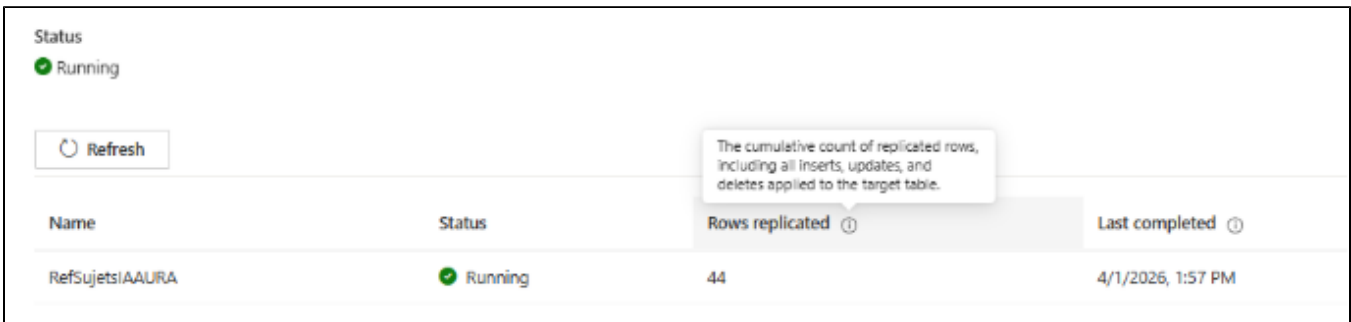
- The underlying Delta table data
- The physical storage (Parquet + metadata)
- Historical versions of the table using Spark

0. Context of the test

- I see 42 rows because I actually have 42 rows in the list (first image)



- however, the replication shows 44 rows because it appears to have stored the changes (second image):
 - "The cumulative count of replicated rows, including all inserts, updates, and deletes applied to the target table."



Is there a way to see those 44 lines (including updates and deletions) instead of the 42?

1. Reading the Delta Table (Current Version)

Even for mirrored tables, the data is stored as a standard Delta table in OneLake and can be accessed via its ABFSS path.

Delta table - data

```
path = "abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA"
df = spark.read.format("delta").load(path)

display(df)
```

This returns the **latest state** of the table, equivalent to what is visible in Fabric after synchronization.

2. Exploring Underlying Storage (Metadata & Files)

The Delta table is physically composed of:

- Parquet data files
- Transaction log (`_delta_log`)
- Additional metadata (indexes, deletion vectors)

You can list these files using:

```
path = "abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA"
files = spark._jvm.org.apache.hadoop.fs.FileSystem \
    .get(spark._jsc.hadoopConfiguration()) \
    .listStatus(spark._jvm.org.apache.hadoop.fs.Path(path))

for f in files:
    print(f.getPath().toString())
```

Example

- [abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA/_delta_log](#)
- [abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA/_index_bin](#)
- [abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA/deletion_vector_e7de9379-f721-4637-8585-2d37ec53ad1b.bin](#)
- [abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA/part-00000-3e4b092f-2246-4b05-a541-10d61f776d3b.c000.zstd.parquet](#)
- [abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA/part-00000-a9a515f8-02ab-48c9-8fb9-9688eaf4567a.c000.zstd.parquet](#)
- [abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA/metadata](#)

Example output:

_delta_log/

_index_bin/

deletion_vector_*.bin

part-*.parquet

metadata/

Key components:

- `_delta_log/` transaction history (versions)
- `part-*.parquet` actual data
- `deletion_vector` logical deletes (CDC optimization)

3. Accessing Historical Versions (Time Travel)

Although not exposed in the Fabric UI for mirrored tables, Delta Lake versioning is still fully available via Spark.

```

from delta.tables import DeltaTable
from pyspark.sql import functions as F

path = "abfss://69300957-941f-4c8a-970f-49b3fcel16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8856-30a30f7cblbf/Tables/dbo/RefSujetsIAAURA"

# 1) Historique Delta
delta_table = DeltaTable.forPath(spark, path)
history_df = delta_table.history() # reverse chronological order

display(history_df)

# 2) Liste des versions disponibles, triées dans l'ordre croissant
versions = [
    row["version"]
    for row in history_df.select("version").distinct().orderBy("version").collect()
]
print("Versions disponibles :", versions)

# 3) Création d'un DataFrame par version
dfs_by_version = {}

for v in versions:
    df_v = (
        spark.read
        .format("delta")
        .option("versionAsOf", v)
        .load(path)
        .withColumn("_delta_version", F.lit(v))
    )
    dfs_by_version[v] = df_v

print(f"{len(dfs_by_version)} DataFrames créés")
print("Exemple version la plus récente :", max(dfs_by_version.keys()))

```

| 12L version | timestamp | ABC userId | ABC userName | ABC operation | ANY operationParameters | ANY job | ANY notebook | ABC clusterId |
|-------------|-------------------------|------------|--------------|---------------|-------------------------|---------|--------------|---------------|
| 3 | 2026-04-01 11:57:55.13 | NULL | NULL | Merge | NULL | NULL | NULL | NULL |
| 2 | 2026-04-01 11:56:09.282 | NULL | NULL | Merge | NULL | NULL | NULL | NULL |
| 1 | 2026-04-01 11:55:55.356 | NULL | NULL | ReplaceTable | NULL | NULL | NULL | NULL |
| 0 | 2026-04-01 11:55:38.943 | NULL | NULL | Write | NULL | NULL | NULL | NULL |

Example

Version 2 :

```
display(dfs_by_version[2])
```

| | | | | | | | |
|----|-----|---------|----------------|------|--------|------|------------|
| 42 | TCI | Syensqo | Mirroring S... | 2026 | Projet | NULL | SyensqoRLS |
|----|-----|---------|----------------|------|--------|------|------------|

Version 3 - Updated made on one field (most recent version so equivalent from delta table synchronization) most recent version

```
display(dfs_by_version[3])
```

42

TCI

Syensqo

Mirroring S...

2026

Projet

Fini

SyensqoRLS

4. Key Observations

- Mirrored tables are **standard Delta tables under the hood**

But UI exposed only the last version (SCD1 type (overwrite))

- All Delta capabilities (time travel, versioning) remain accessible via Spark
- The Fabric UI does **not expose version history** for mirrored tables
- Each version represents a **full snapshot**, not just incremental changes, no change lo apparently
- CDC changes are internally managed through:
 - Delta commits
 - Deletion vectors
 - Transaction logs

It seems we can access to commit logs

The screenshot shows a Databricks notebook interface. At the top, there is a code cell with the following Spark SQL commands:

```
1 log_path = "abfss://69300957-941f-4c8a-970f-49b3fce16e0d@onelake.dfs.fabric.microsoft.com/f73eb213-d657-4e43-8956-30a30f7cb1bf/Tables/dbo/RefSujetsIAAURA/_de
2
3 raw_log_df = spark.read.json(log_path)
4
5 print("Colonnes du log :")
6 print(raw_log_df.columns)
7
8 display(raw_log_df)
```

Below the code cell, the execution status is shown: "2 sec - Command executed in 2 sec 504 ms by Théo COLOMBANI on 01/04/2026 16:43:56".

The notebook output shows the columns of the log: ["add", "commitInfo", "metaData", "protocol", "remove"].

Below the output, there is a table view of the data. The table has 12 rows and 5 columns: "add", "commitInfo", "metaData", "protocol", and "remove". The "add" column contains NULL values. The "commitInfo" column contains JSON strings. The "metaData" column contains NULL values. The "protocol" column contains JSON strings. The "remove" column contains NULL values.

| # | add | commitInfo | metaData | protocol | remove |
|----|-----------------|---|--------------------|---|-------------------|
| 1 | NULL | NULL | NULL | {"minReaderVersion":3,"minWriterVersion":7,"readerFeatures":... | NULL |
| 2 | NULL | NULL | {"format":{"pro... | NULL | NULL |
| 3 | {"stats":{"v... | NULL | NULL | NULL | NULL |
| 4 | NULL | {"additionalCommitInfo":{"LastVacuumTimestamp":"4/1/20... | NULL | NULL | NULL |
| 5 | NULL | NULL | NULL | {"minReaderVersion":3,"minWriterVersion":7,"readerFeatures":... | NULL |
| 6 | NULL | NULL | {"format":{"pro... | NULL | NULL |
| 7 | NULL | {"additionalCommitInfo":{"LastVacuumTimestamp":"4/1/20... | NULL | NULL | NULL |
| 8 | {"stats":{"v... | NULL | NULL | NULL | NULL |
| 9 | {"stats":{"v... | NULL | NULL | NULL | NULL |
| 10 | NULL | NULL | NULL | NULL | ["path":{"part... |
| 11 | NULL | {"additionalCommitInfo":{"LastVacuumTimestamp":"4/1/20... | NULL | NULL | NULL |
| 12 | NULL | {"additionalCommitInfo":{"LastVacuumTimestamp":"4/1/20... | NULL | NULL | NULL |

But without testing more tok now which row was deleted or added

```
1 remove_df = raw_log_df.filter("remove is not null")
2 display(remove_df.select("remove.path", "remove.deletionTimestamp", "remove.dataChange"))
```

[5] ✓ 1 sec - Command executed in 1 sec 308 ms by Théo COLOMBANI on 01/04/2026 16:44:15

> Spark jobs (2 of 2 succeeded) Resources

Table view Download

| # | src path | % deletionTimestamp | % dataChange |
|---|---------------|---------------------|--------------|
| 1 | part-00000... | 1775044674950 | true |

```
1 add_df = raw_log_df.filter("add is not null")
2 display(add_df.select("add.path", "add.size", "add.modificationTime", "add.dataChange"))
```

[6] ✓ 1 sec - Command executed in 1 sec 2 ms by Théo COLOMBANI on 01/04/2026 16:44:28

> Spark jobs (2 of 2 succeeded) Resources

Table view Download

| # | src path | % size | % modificationTime | % dataChange |
|---|---------------|--------|--------------------|--------------|
| 1 | part-00000... | 20595 | 1775044554839 | true |
| 2 | part-00000... | 20595 | 1775044674950 | true |
| 3 | part-00000... | 6980 | 1775044674387 | true |