



Optimizing Data Files in Apache Iceberg: Performance strategies - Q&A

Querying 100s of petabytes of data demands optimized query speed specifically when data accumulates over time. We have to ensure that the queries remain efficient because over time you may end up with a lot of small files and your data might not be optimally organized.

In this talk, we will cover:

- Apache Iceberg table format
- Problems in the data lake: small files, unorganized files
- Techniques such as: partitioning, compaction, metrics filtering
- Overlapping metrics problem
- Solving it using sorting, Z-order clustering



Alex Merced
Developer Advocate
Dremio



[Dipankar Mazumdar](#)
Developer Advocate
Dremio

Webinar Q&A

1. Is Iceberg format similar to Delta format?

Answer: It is similar in the purpose - i.e. they are both table formats on top of data lakes, which helps you manage those raw data files like Parquet, ORC, etc. and allows you to do things like ACID transactions, time-travel, version rollbacks, etc.

In terms of the features, there are some differences between both these formats. You can read more here:

<https://www.dremio.com/blog/comparison-of-data-lake-table-formats-apache-iceberg-a-pache-hudi-and-delta-lake/>

2. Is Iceberg table format a virtual/lite layer on top of the actual data files?

Answer: Apache Iceberg is a layer on top of the actual data files, yes. They provide a set of APIs and libraries to interact with these data files.

3. Can you please go over hidden partitioning again compared to hive ?

Answer: Here are some more detailed read:

- <https://iceberg.apache.org/docs/latest/partitioning/>
- <https://www.dremio.com/blog/fewer-accidental-full-table-scans-brought-to-you-by-apache-icebergs-hidden-partitioning/>

4. Does rewrite_data_files also regenerate manifests?

Answer: rewrite_data_files procedure is used to compact data files or sort the data using normal sorting or Z-ordering. Details:

https://iceberg.apache.org/docs/latest/spark-procedures/#rewrite_data_files

5. While running Z-Ordering, will there be any slowdown on inserts/updates

Answer: In general, executing the rewrite_data_files procedure with Z-order as a sort strategy should not impact other operations but there should be certain considerations.

6. Data Warehouse has transactions so that parallel writes can properly happen, how does iceberg properly handle this?

Answer: Iceberg supports multiple concurrent writes using Optimistic concurrency control. Each writer assumes that no other writers are operating and writes out new table metadata for an operation. Then, the writer attempts to commit by atomically swapping the new table metadata file for the existing metadata file. If the atomic swap fails because another writer has committed, the failed writer retries by writing a new metadata tree based on the new current table state.