



The Who, What and Why of Data Lakehouse Table Formats - Q&A

In the world of data acceleration and optimization, both materialized views and Dremio's Data Reflections stand out as pivotal tools. Our upcoming webinar titled "Materialized Views and Dremio's Reflections" aims to demystify these technologies, comparing their benefits, limitations, and unique features. Dive deep into understanding the core differences between materialized views and Dremio's Reflections. Whether you're a seasoned data professional or just starting out, this webinar offers insights to optimize your data strategy. Discover the nuances, best practices, and real-world applications of these powerful tools, and make informed decisions for your data architecture.



Alex Merced

Developer Advocate
Dremio

Show's Q&A

1. Refusing reflection—who is authorized to trigger reflection or it can be set up to be automatically done?

Answer: Admin users and users who were granted the privilege to create reflections can create reflections. End consumers will often not be aware the reflections exist and just enjoy better performance.

2. How does the use say: only use a reflection which is uptodate to the time the query is running?

Answer: Reflection are Iceberg tables so if a query is executed while the reflection is updating it would use the previous reflection snapshot or not use the reflection at all depending on the query.

3. Are materialized views incremental refreshed, particularly when materialized views have joins in them?

Answer: if the underlying data is in an iceberg table, anytime the table is updated any reflections based on the data or downstream views will update, so if you have reflections on a join of two iceberg tables when either table updates a reflection on a join will update

4. Can Reflections be created on External DBs other than Data lake?

Answer: Yes, any data source connected to Dremio can have reflections enabled, only Iceberg tables are updated incrementally. Other sources will get a full refresh.

5. Is it better for the user to create reflection (Static) or allow Dremio to determine the creation of reflection based on the query patterns (Dynamic). Is this option available in cloud and on-prem installations?

Answer: Dremio reflection recommendations is meant to make it easier to decide which reflections will to create. Dremio will continue to expand the ways it aims to make reflections a feature that just takes a lot of the work in having a performant lakehouse.

6. How does this interact with dbt? Can we build the data reflections there?

Answer: I think you can run any valid DremioSQL with the dbt integration, and reflections can be created with SQL

7. Do we have a scheduler to refresh reflection?

Answer: You can set a frequency for reflection refreshes

8. How to Refresh a Reflection?

Answer: you can either manually trigger. Refresh with SQL or set a schedule via SQL

9. Do we have to create multiple reflection based on access pattern?

Answer: you don't have to, ideally you only create reflections when necessary. If queries without reflections meet you SLAs then just let Dremio run the query but if you need additional performance then reflections offers a way to get there.