

YSQL Materialized Views

Fizaa Luthra
Friday, May/17/2022

YFTT

YUGABYTEDB
FRIDAY
TECH TALKS



yugabyte**DB**

Overview

- Materialized views are relations that contain the result of a query.
- Unlike regular views, they persist the results of the view query.
- Unlike regular tables, they cannot be directly updated. Instead, we must use the REFRESH command to update the MV.
- Used to speed up access to the result of a query (at the expense of some storage).
- Currently YSQL only (no MVs in YCQL).

Materialized View Commands

- CREATE MATERIALIZED VIEW
- REFRESH MATERIALIZED VIEW
- DROP MATERIALIZED VIEW

CREATE MATERIALIZED VIEW

- Create the matview as a YB relation
- Can specify the query, column names
- Can specify WITH DATA (default) / WITH NO DATA

Example:

```
CREATE MATERIALIZED VIEW mv AS SELECT COUNT(*) FROM  
base_relation GROUP BY col_name;
```

REFRESH MATERIALIZED VIEW

- Used to update the contents of the matview
- Can specify WITH DATA (default) / WITH NO DATA
 - WITH DATA: re-execute the query and re-populate the matview
 - WITH NO DATA: mark the matview as unpopulated (or “outdated”), leaving it in an unscannable state.
- Two modes of REFRESH WITH DATA:
 - Default option
 - “CONCURRENTLY” option

REFRESH MATERIALIZED VIEW (default option)

- Rebuilds and reindexes the relation
- Better if there are a large number of new/modified rows

Example:

```
REFRESH MATERIALIZED VIEW mv;
```

Note: Does not block concurrent SELECTs

REFRESH MATERIALIZED VIEW (CONCURRENTLY)

- Construct a “diff” relation
- Perform updates accordingly
- Better if there are a fewer number of new/modified rows
- Requires a unique index on the materialized view

Example:

```
REFRESH MATERIALIZED VIEW CONCURRENTLY mv;
```

DROP MATERIALIZED VIEW

- Used to DROP the materialized view
- Can specify CASCADE / RESTRICT

Example:

```
DROP MATERIALIZED VIEW mv;
```


Regular Views vs. Materialized Views

Regular Views

- Do not persist the results of a query. The query is run every time we reference the view.
- Don't require any maintenance.
- Generally used when the data is accessed infrequently.

Materialized Views

- Persist the results of a query. The query is only re-run when the `REFRESH` command is used.
- Require maintenance (a `REFRESH` must be performed to retrieve the updated result).
- Used to speed up the access to results of a query (at the expense of some storage).
- Behave like a cache for the view query.
- Generally used when the data is accessed frequently.

Thank You

Join us on Slack: yugabyte.com/slack (#yftt channel)

Star us on Github: github.com/yugabyte/yugabyte-db

