dremio

Dremio Software

Optimize Metadata Refresh Frequency

Introduction

This document provides best practices for setting and adjusting the metadata frequencies for datasets in Dremio. It introduces and provides an overview of metadata in general and describes in detail how to best manage it in Dremio.

What is Metadata

Metadata is information about a dataset, including high-level information like dataset name, schema, and partition layout, and low-level details like statistics, locality, shards, file information for file system sources, etc.

The Importance of Refreshing Metadata

When you submit a query to Dremio, Dremio uses metadata to validate the query and generate an accurate and efficient query plan. When you initially promote a dataset, Dremio automatically collects metadata for the dataset.

For file system data sources like S3, HDFS, ADLS, or GCP, if a dataset's metadata is not up to date, Dremio queries cannot return the latest data. Even though new data may be available in the data lake, Dremio is unaware of it if metadata has not been refreshed. Stale or out-of-date metadata can result in inefficient query plans for non-file system sources.

Inline Metadata Refresh

If you submit a query against a table with expired metadata, Dremio performs an inline metadata refresh. In this case, the submitted query pauses until the metadata refresh is complete. Inline metadata refresh can add considerable time to a query for large datasets and external datasets. To avoid extended query times due to expired metadata, you can refresh metadata before it expires.

> ⚠ CAUTION
> When configuring metadata refresh in the Dremio console, ensure that the expiration frequency is larger than the refresh frequency.

Metadata for Iceberg Tables

You don't need to refresh metadata for Iceberg tables. When data is added or updated in an Iceberg table, the table's metadata updates with the DML operation, so no separate refresh is needed. Outside of DML capabilities, this is another excellent feature of using Iceberg tables -- your data is always up to date, with one specific exception noted below.

If all or most of your datasets are Iceberg tables, there is no need to set up an isolated engine for metadata refresh as it can be underutilized.

> ℹ️ NOTE
> If your catalog type is HDFS and you are performing DML operations outside of Dremio using another engine like Spark, then you need to explicitly refresh the Iceberg table metadata

> using the ALTER TABLE command. ALTER TABLE is a lightweight operation performed only on the coordinator to store the latest snapshot file location.

Metadata for Delta Lake Tables

Metadata refresh is required for Delta Lake tables and is performed on the coordinator. For Delta Lake tables, metadata is read from the delta log and is a lightweight operation.

Expiring Metadata

For non-Iceberg tables, when new files are added to the data lake, Dremio does not recognize the new files until you refresh the metadata. Dremio executes queries based on the current metadata.

Set the frequency for metadata expiration according to your requirements, ensuring that users are not submitting queries based on stale metadata. If metadata expires before it is refreshed, Dremio performs an inline refresh, which can add latency to query execution times.

Metadata Refresh Recommendations

Adjust Refresh Frequency

Dremio recommends reviewing the metadata refresh frequencies for all your data sources every quarter at a minimum, or any time you add a new data source. This scheduled review ensures you can keep your refresh frequencies configured according to how often metadata changes in the data source.

Metadata refresh settings should align with the frequency of changes or ingestion in your data sources and the trade-off between compute resources and resources needed for metadata refresh. If you are ingesting new data frequently, you will need to decide on an acceptable period before metadata is refreshed. If you are ingesting data once a day, you only need to refresh metadata once a day. In this case, a higher refresh frequency will provide no benefit; it will consume compute resources that are necessary for query execution and add unnecessarily to operating costs.

Scheduled Refresh

The Dremio console provides metadata refresh settings at the data source level. The default metadata refresh for new data sources is every hour. For most data sources, this is far too frequent. For example, if the data in a source is only updated once every six hours, it does not make sense to refresh metadata hourly. You should update the refresh schedule to match the frequency of data updates in your source. See Scheduling Metadata Refreshes in the Dremio docs.

If you have multiple datasets with different update frequencies at the source level, consider creating an independent data source for each group of datasets to align metadata refresh frequencies with ingestion frequency better.

> ℹ️ NOTE
> Dremio does not support scheduled metadata refresh at the table level.

On-Demand Refresh via PDS Override

Since metadata refresh can be scheduled at the data source level and overridden programmatically for individual tables, you should review each new data source to understand the best refresh method. For example, on data lake sources you might set a very long metadata refresh schedule on the source to ensure that scheduled refresh doesn't happen. Still, when data generation is complete, you can perform an `ALTER TABLE .. REFRESH METADATA` as part of the ETL process.

It might make sense to set a more extended refresh schedule for relational sources at the source level, but override those source settings on tables because they will be updated more frequently.

It only makes sense to refresh the metadata for datasets that are updated during overnight ETL runs once you know the ETL process is finished. In these cases, you can create a script that triggers the manual refresh of each dataset at the end of the ETL process. The script can call the SQL REST API or run JDBC/ODBC queries to execute the `ALTER TABLE .. REFRESH METADATA` command.

Suppose the ETL process does not fully update an existing dataset but only changes partitions or creates new partitions. In that case, you can use the script to tell Dremio to refresh only the changed or new partitions -- this works only with Parquet and Iceberg datasets, not CSV or JSON.

Scheduled refresh does not make sense for data sources that include many datasets where only a small number will change structure or have new files added. In this case, set the metadata to never refresh at the source and add scripts that trigger a manual refresh on a specific dataset with `ALTER TABLE .. REFRESH METADATA`.

If you use on-demand refresh, be sure that you set metadata at the data source level to never refresh and never expire. You can then trigger an on-demand refresh for all datasets that you update.

> ℹ️ NOTE
> If you don't schedule the metadata refresh and you don't refresh with a script, queries against stale or invalid metadata will trigger an inline refresh during the planning phase. Inline refresh can increase the duration of query execution.

Use a Dedicated Engine

Dremio recommends using a dedicated engine for metadata refresh if executor nodes are processing many metadata refresh jobs. Using a dedicated engine ensures that all metadata refresh activities on the executor are isolated from other workloads and keeps CPU and memory resources dedicated to business-critical workloads. For file-based sources with a high number of files (i.e., more than 1 million), it is better to use multiple small engine instances than a single scale-up instance because multiple instances can run the metadata refresh queries in parallel.

Use the `query_type()` engine routing rule for metadata refresh jobs (`query_type() = 'Metadata Refresh'`).