Data Lake Analytics Special Interest Group

SIG-DLA covers the development of StarRocks Data Lake. To know the background and goals, please see here.

SIG-DLA Leaders: TBD

Contact: Slack, search channel #sig-dla
Important Links:

- GitHub: https://github.com/StarRocks/starrocks
- StarRocks Iceberg V2 design doc
- Open Issues:
    - https://github.com/StarRocks/starrocks/issues/5062
    - https://github.com/StarRocks/starrocks/issues/7268

If you have topics for the SIG meetings, feel free to add them to the upcoming meeting dates.

Agenda / Minutes

July 22nd, 2022

Time: 16:30 CST; 1:30 PDT
Host: Kate Shao

Topics:

- [Kang Zhou] StarRocks Iceberg V2 design doc
    - Project background: Currently, StarRocks supports Apache Iceberg table spec V1, but as more and more users use Iceberg, the demand for StarRocks to support Apache Iceberg table spec V2 is increasing. Therefore, we'd like to add Apache Iceberg V2 support in the next version. The main difference between V2 and V1 is that V2 supports a **row-level deletion format,** for more details, see: https://iceberg.apache.org/spec/

- Three proposed solutions:

Option 1: JNI

1. **JNI**
    - Implementing StarRocksDeleteFilter
    - Implementing StarRocks Row Wrapper
2. **FE**
    - When creating the FileScanTask, get the DeleteFile information and send it to BE
3. **BE**
    - Initialize DeleteFilter
    - When reading chunks, refine the schema to be read according to the information provided by DeleteFilter
    - For each chunk, JNI calls StarRocksDeleteFilter for filtering
4. **Future optimization**
    - Use off-heap to swap data and avoid memory copies

Cons

- Involves multiple memory copies and performance problems
- The native parquet/orc reader is not used, so the relevant optimization cannot be used

Option 2: C++ implementation

1. **FE**
    - When creating the FileScanTask, get the DeleteFile information and send it to BE
2. **BE**
    - Implement the logic of DeleteFilter
    - Implement the logic of PositionStreamDeleteFilter
    - Implement the logic of PositionSetDeleteFilter
    - Implement Predicate related logic
    - Implement Avro read-related logic
3. **Future optimization**
    - Performance
        - Rewriting workload
        - The subsequent synchronization of iceberg's SDK modification also requires work

Option 3: Rewrite the iceberg SDK

- It requires collaboration with the Iceberg community to implement a C++ or rust SDK. It might take a long time.

Discussions:

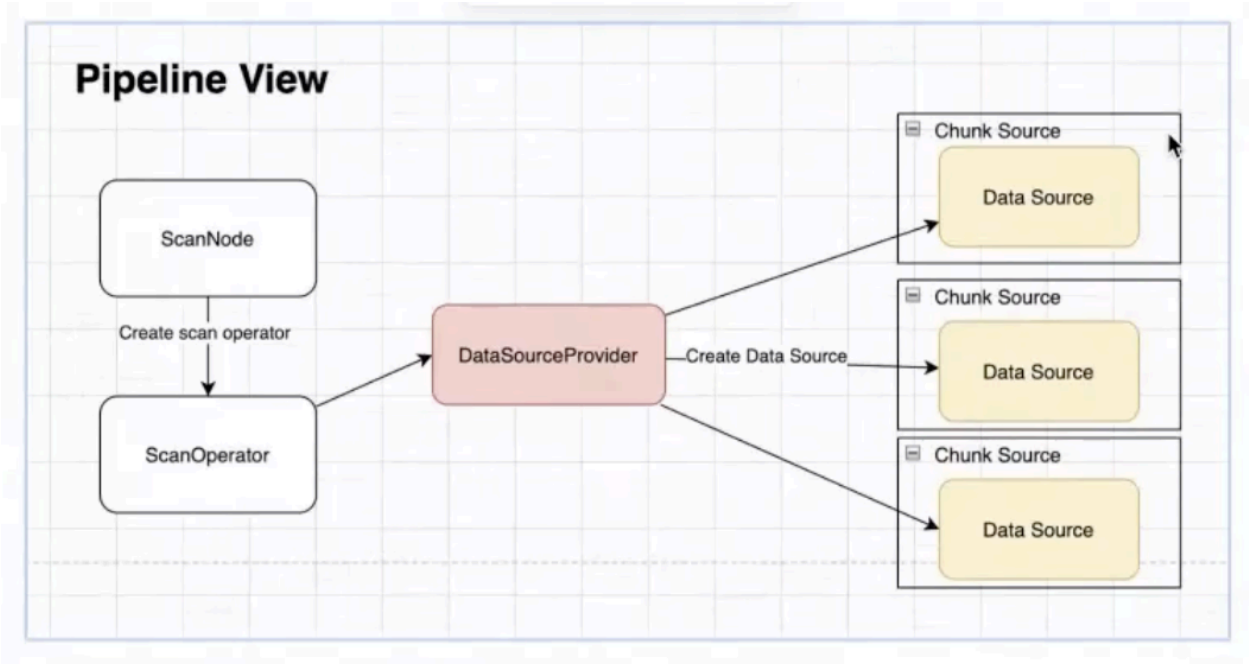
- 【Choury】Option 2 might be the easiest to do, so I think it'd be better to go with option 2.

June 20th, 2022

Time: 16:30 CST; 1:30 PDT
Host: Dorian Zheng

Topics:

- [Stephen Li] Introduction to the new DLA Framework
    - Why we started the new connector design
    - Goals- StarRocks to connect to external data sources more easily
    - Catalog (internal and external catalog), architecture and FE design overview
    - Roadmap
        - Phase 1
        - Phase 2
    - Hive connector preview in StarRocks 2.3. Scheduled to release end of June.
- [Kang Zhou] The status on the BE side:
    - Work done: defined ScanNode and ScanOperator; MySQL, JDBC, ES, HDFS data sources already adapted to the framework
    - Next step: implementation of communication protocol between FE & BE



Discussions:

[Heng Zhao] Plan for phase 2 sync function?
[liubang]: Plan for Iceberg writer? Do we need to write a C++ SDK?