



What's New in Apache Iceberg?




Gnarly Data Waves • June 6th 2023

About Me



- **Current:** Developer Advocate @Dremio
- **Open source:** Apache Iceberg, Arrow, Project Nessie
- **Past:** BI, Machine Learning, Data architecture



Agenda

1 Apache Iceberg Community

2 Apache Iceberg 1.2.0 Release

3 New Capabilities

4 Pylceberg: Pythonic API

5 Compute Engine Updates

Apache Iceberg Community

“A developer community is not just about code; it's diversity & about creating a space where ideas flourish, and individuals inspire each other”




Apache Iceberg Workspace, Developer Mailing Lists




@IcebergDevs


Apache Iceberg 1.2.0: Core Library



Branching, Tagging



Delta to Iceberg migration



Human-readable metrics

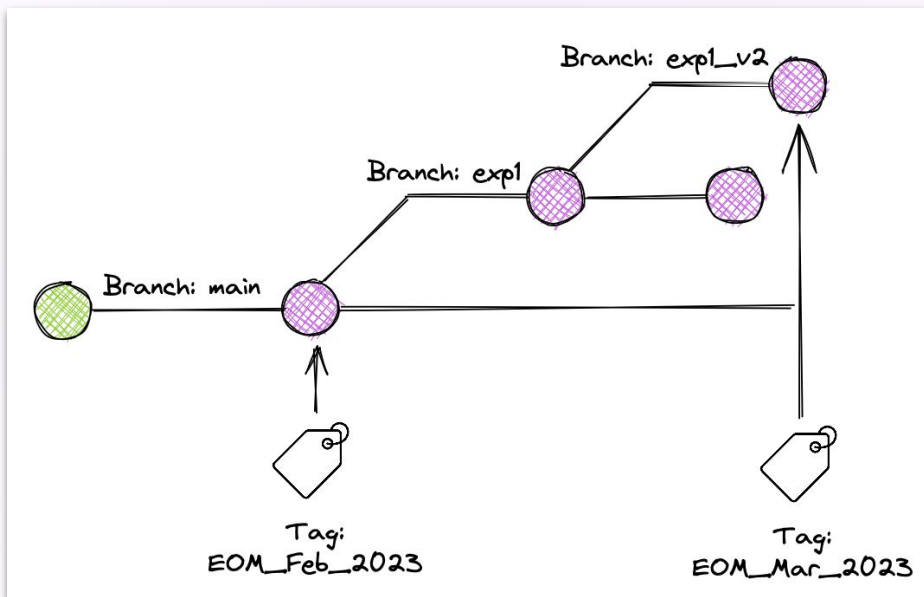


Stats File reusability



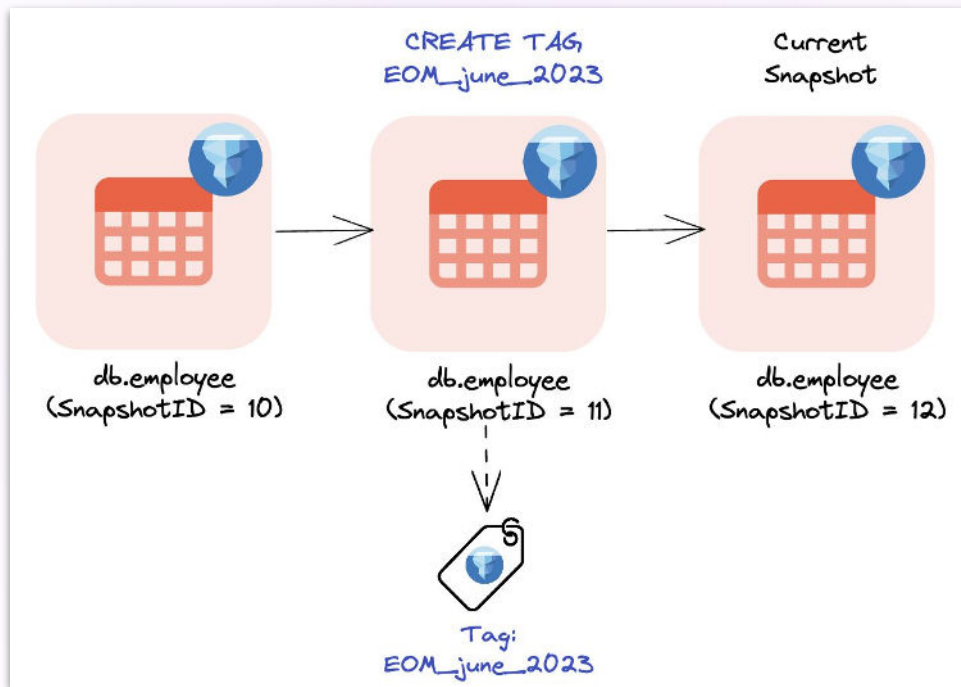
FileIO support for ORC readers & writers

Branching & Tagging



- Git-like Branch & Tags
- Keep track of the state[version] of Iceberg tables
- Core: Snapshot

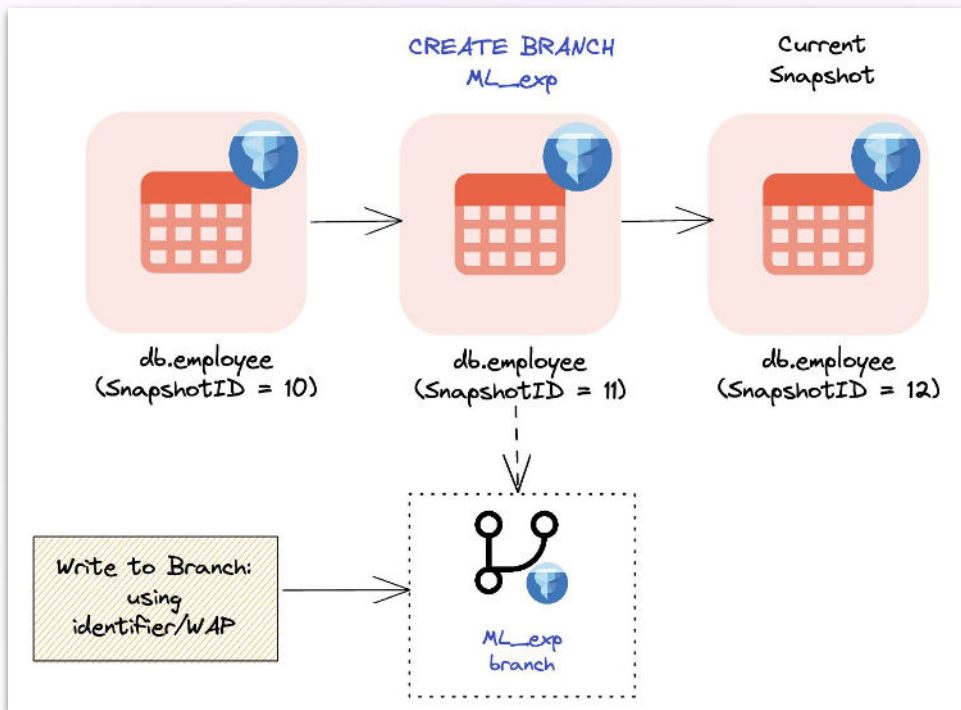
Iceberg Tags



- Immutable **labels** for individual snapshots
- RETAIN tags
- Data/ML reproducibility, Version-specific analysis

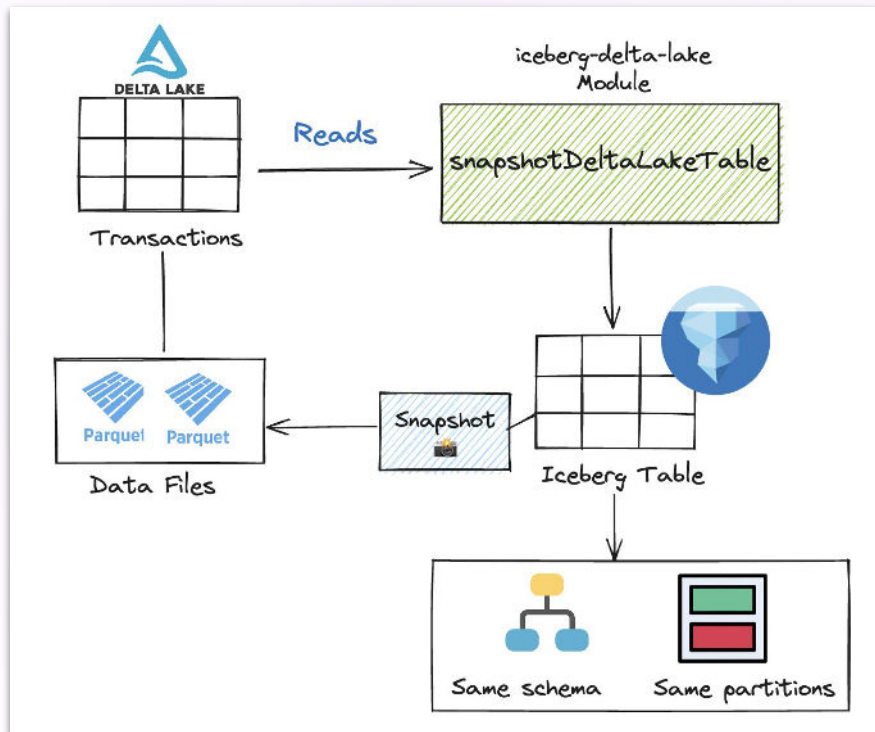
```
ALTER TABLE glue.test.employee  
CREATE TAG EOM_June_2023
```


Iceberg Branches



- Mutable **labels** for snapshots
- Adheres to Iceberg's commit conflict resolution & retry policy
- Write to a Branch using **identifier** or **WAP**
- Write-Audit-Publish, Isolated data experiments

Delta to Iceberg migration



- Reads Delta Lake table transactions -> Converts to Iceberg
- Schema & Partitioning scheme maintained
- Existing data files added to Iceberg table's metadata
- Snapshot uses original table's data files

Human-readable metrics

BEFORE

```
"lower_bound": {  
  "bytes": "1\u0006\u0000"  
},  
  
"upper_bound": {  
  "bytes": "1\u0006\u0000"  
}
```

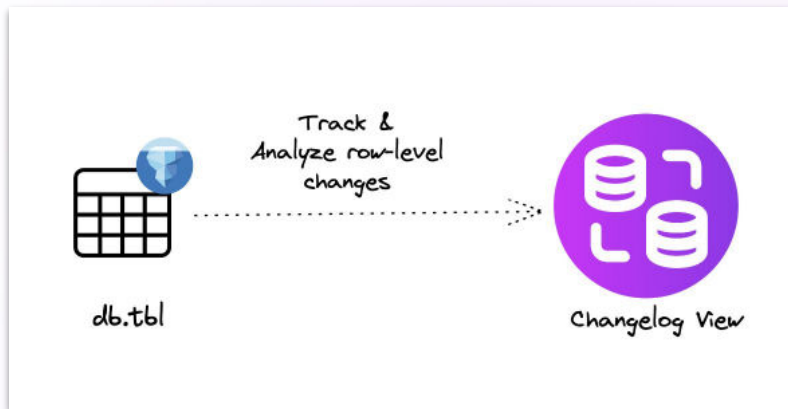
Files Metadata
Table

AFTER

```
SELECT readable_metrics.id.lower_bound,  
readable_metrics.id.upper_bound FROM  
test.employee.files"
```

	readable_metrics.id.lower_bound	readable_metrics.id.upper_bound
0	6	6
1	7	7
2	8	8
3	4	4
4	3	3
5	2	2
6	1	1

Change Data Capture



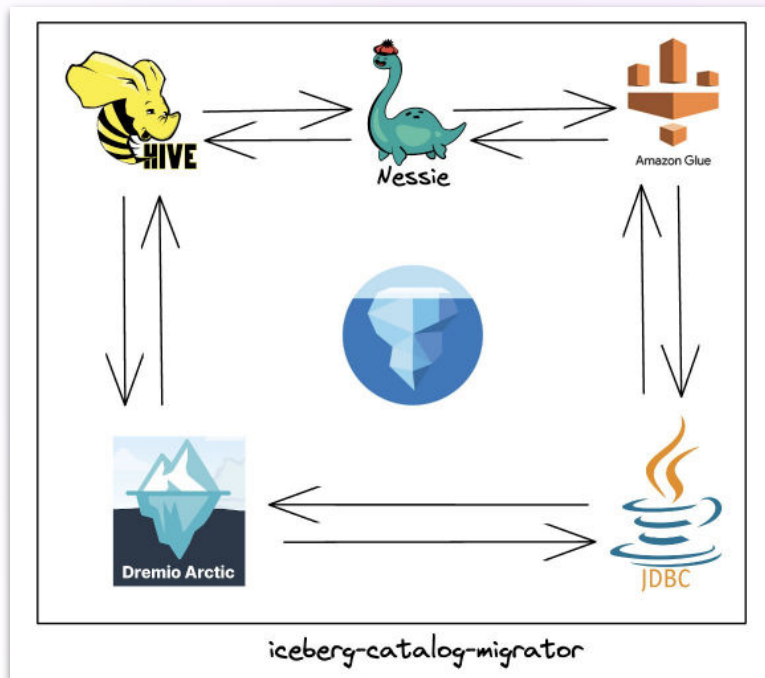
Spark Procedure

```
CALL spark_catalog.system.create_changelog_view(  
  table => 'db.tbl',  
  options => map('start-snapshot-id', '1', 'end-snapshot-id', '2')  
)
```

```
SELECT * FROM glue.db.tbl.changes
```

id	name	_change_type	_change_ordinal	_change_snapshot_id
1	Alice	INSERT	0	5390529835796506035
2	Bob	INSERT	0	5390529835796506035
1	Alice	DELETE	1	8764748981452218370

Catalog Migrator Tool

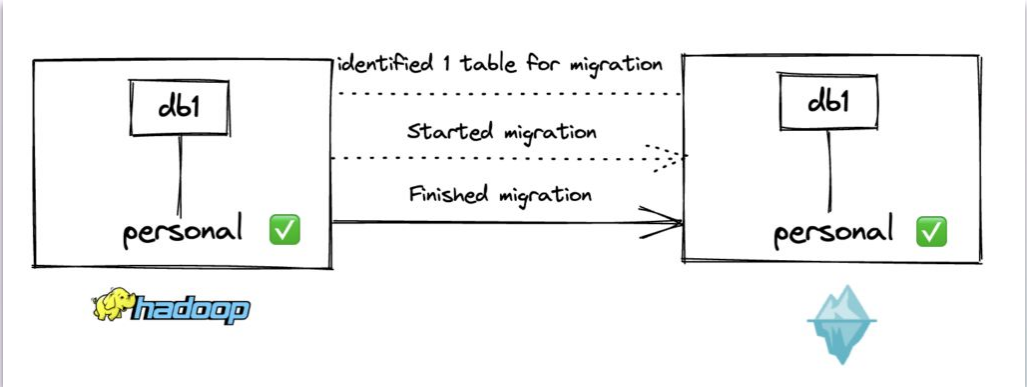


- Bulk migrate from one catalog to another without data copies
- Commands: register & migrate
- CLI tool:
<https://github.com/projectnessie/iceberg-catalog-migrator>

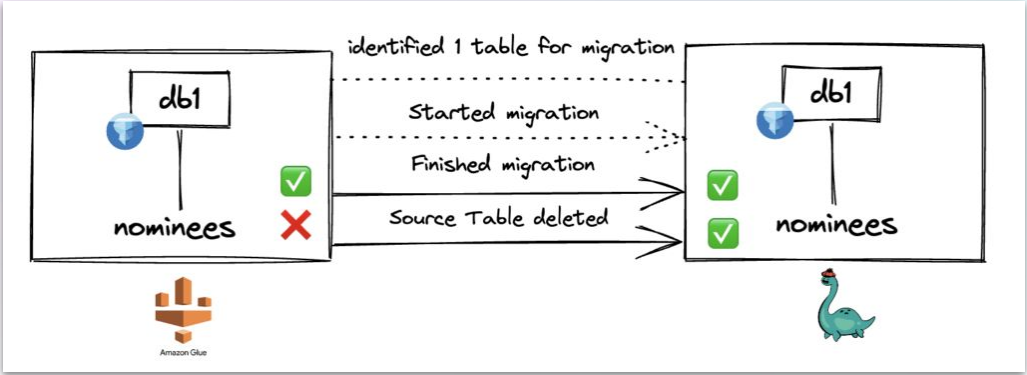
```
Untitled-1

java -jar iceberg-catalog-migrator-cli-0.2.0.jar migrate \
--source-catalog-type GLUE \
--source-catalog-properties warehouse=s3a://bucket/gluecatalog/,io-
impl=org.apache.iceberg.aws.s3.S3FileIO \
--target-catalog-type NESSIE \
--target-catalog-properties
uri=http://localhost:19120/api/v1,ref=main,warehouse=s3a://bucket/nessie/,io-
impl=org.apache.iceberg.aws.s3.S3FileIO \
--identifiers db1.nominees
```

Catalog Migrator Tool




Scenario 1



Scenario 2

Spark: Performance Updates



Storage Partitioned Joins

[PR](#)



Arrow Environment Settings

[PR](#)



Aggregate Pushdown Support

[PR](#)



Default Write
Distribution mode

[PR](#)




Metadata-only DELETE

[PR](#)


Flink Updates




Metadata Tables
[PR](#)



Branching & Tagging
[PR](#)




Throttling Support for Streaming Reads
[PR](#)



Enable Multiple Flink Sinks
[PR](#)

Dremio Sonar Updates



COPY INTO

```
1 COPY INTO context.MyTable
2   FROM '@SOURCE/bucket/path/folder/'
3   FILE_FORMAT 'json'
4
```

```
1 COPY INTO context.myTable |
2   FROM '@SOURCE/bucket/path/folder'
3   REGEX '.*.csv'
4
```

Dremio Sonar Updates



Rollback Iceberg Tables

```
1 ROLLBACK TABLE demo.example_table TO SNAPSHOT '2489484212521283189'  
2  
3  
4  
5
```

```
1 ROLLBACK TABLE demo.example_table  
2 TO TIMESTAMP '2022-06-22 17:06:00'  
3
```

Dremio Sonar Updates




Optimize Iceberg Tables

```
1 OPTIMIZE TABLE demo.example_table
2 |   REWRITE DATA (TARGET_FILE_SIZE_MB=256)
3 |
```

```
1 OPTIMIZE TABLE demo.example_table
2 |   FOR PARTITIONS sales_year=2023
3 |
```

```
1 OPTIMIZE TABLE demo.example_table
2 |   REWRITE MANIFESTS
```

Dremio Sonar



Expire Snapshots:
VACUUM

Context:

```
1 VACUUM TABLE 's3.SF_Incidents2016'  
2   EXPIRE SNAPSHOTS older_than=TIMESTAMP '2023-04-20 00:00:00.000' retain_last=20;  
3
```

PyIceberg: Stats Filtering (available)



Skip Full Data Files: Pruning

```
from pyiceberg.catalog import load_catalog
from pyiceberg.expressions import EqualTo
catalog = load_catalog("glue", **{"type": "glue"})
tbl = catalog.load_table('test.employee')
sc = tbl.scan(row_filter=EqualTo("role", "Developer Advocate")).to_arrow()
sc = sc.to_pandas()
sc
```

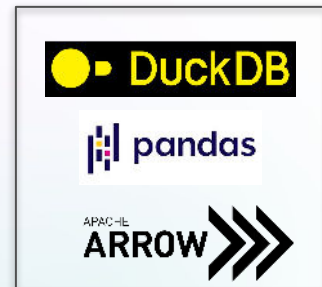

operator lib

skip files based on value

	id	name	role	salary
0	1	Dipankar	Developer Advocate	25000.0

- Skip Full Data Files
- Helps with Pruning

Pylceberg: Convert to Ray



```
ray_dataset = table.scan().to_ray()
```

For Directly converting Iceberg datasets to Ray: ML tasks

PyIceberg: LIMIT

```
from pyiceberg.catalog import load_catalog

table = catalog.load_table("nyc.taxis")
scan = table.scan(
    row_filter=GreaterThanOrEqual("trip_distance", 10.0),
    selected_fields=("VendorID", "tpep_pickup_datetime", "tpep_dropoff_datetime"),
    limit=100,
)
```

Limit preview of an Iceberg table

Pylceberg: Positional Deletes

```
DELETE FROM iceberg_table WHERE id = '10'  
  
select count(uid) from "iceberg_table"  
where id = '10'
```

Fix for deleted records returned back in Pylceberg

Pylceberg CLI: GCS Support



Support for GCS in Pylceberg CLI

Pylceberg: What can be done?



Smooth & Mature integration with other tools

Apache Iceberg 1.3.0

Add support for Spark 3.4 #7174

 Closed

aokolnychiy opened this issue on Mar 22 · 3 comments · Fixed by #7378

Spark 3.4: Add RewritePositionDeleteFilesProcedure #7572

 Merged

aokolnychiy merged 3 commits into `apache:master` from `szehon-ho:rewrite_deletes_procedures` 2 weeks ago

Improve distribution and ordering in Spark 3.4 #7633

 Closed

aokolnychiy opened this issue 2 weeks ago · 0 comments · Fixed by #7637

Spark 3.4: Avoid local sort for MERGE cardinality check #7558

 Merged

aokolnychiy merged 2 commits into `apache:master` from `aokolnychiy:avoid-sort-merge-cardinality-check` 2 weeks ago

Thank you!