



Making Iceberg walk and talk like Postgres

Marco Slot

Principal Software Engineer, Snowflake

DuckCon • June 24, 2026



pg_lake: Iceberg in Postgres

pg_lake is a Postgres extension that adds:

1. Iceberg tables in Postgres
2. Query/import/export files in object storage

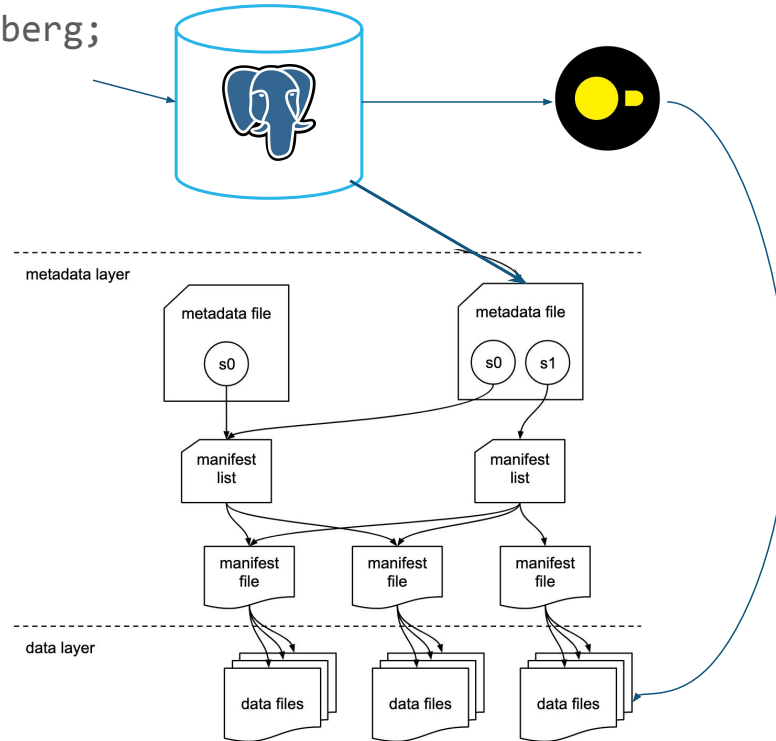
Bottomless storage,
Fast analytical queries,
Transactional data pipelines.

https://github.com/Snowflake-Labs/pg_lake

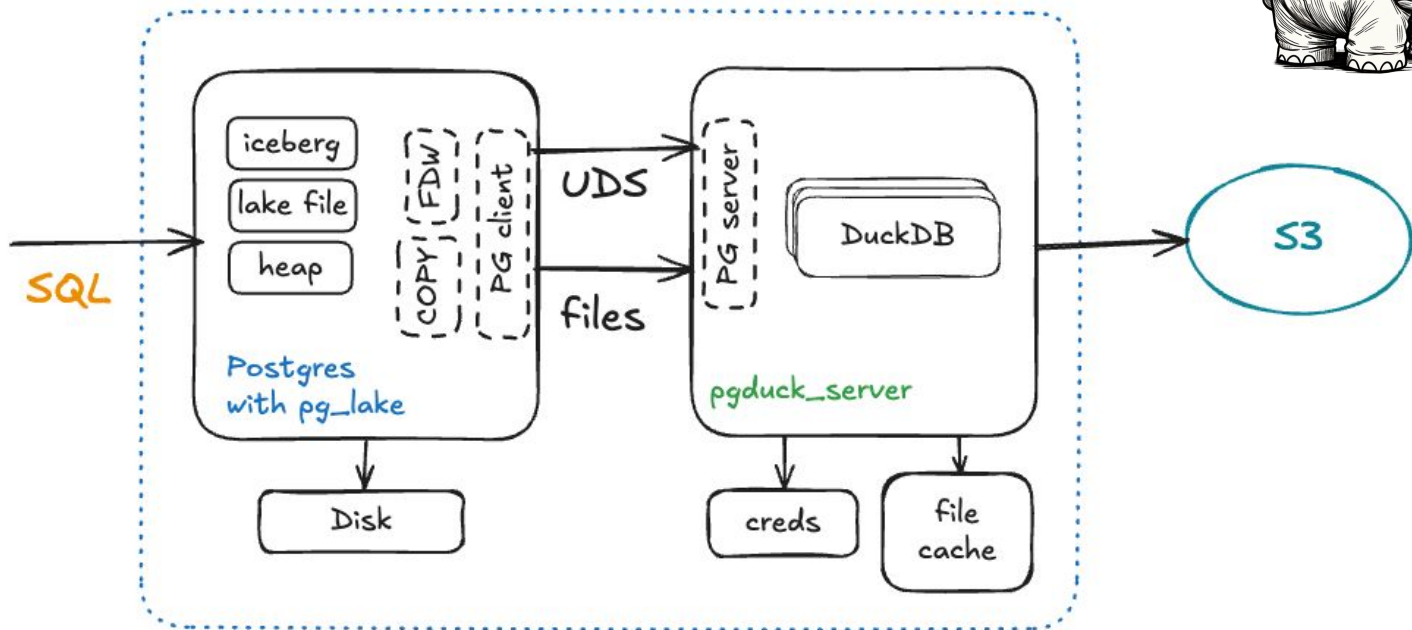



Iceberg tables in Postgres

```
CREATE TABLE pgice (...) USING iceberg;  
INSERT INTO pgice ...;  
SELECT .. FROM pgice;
```

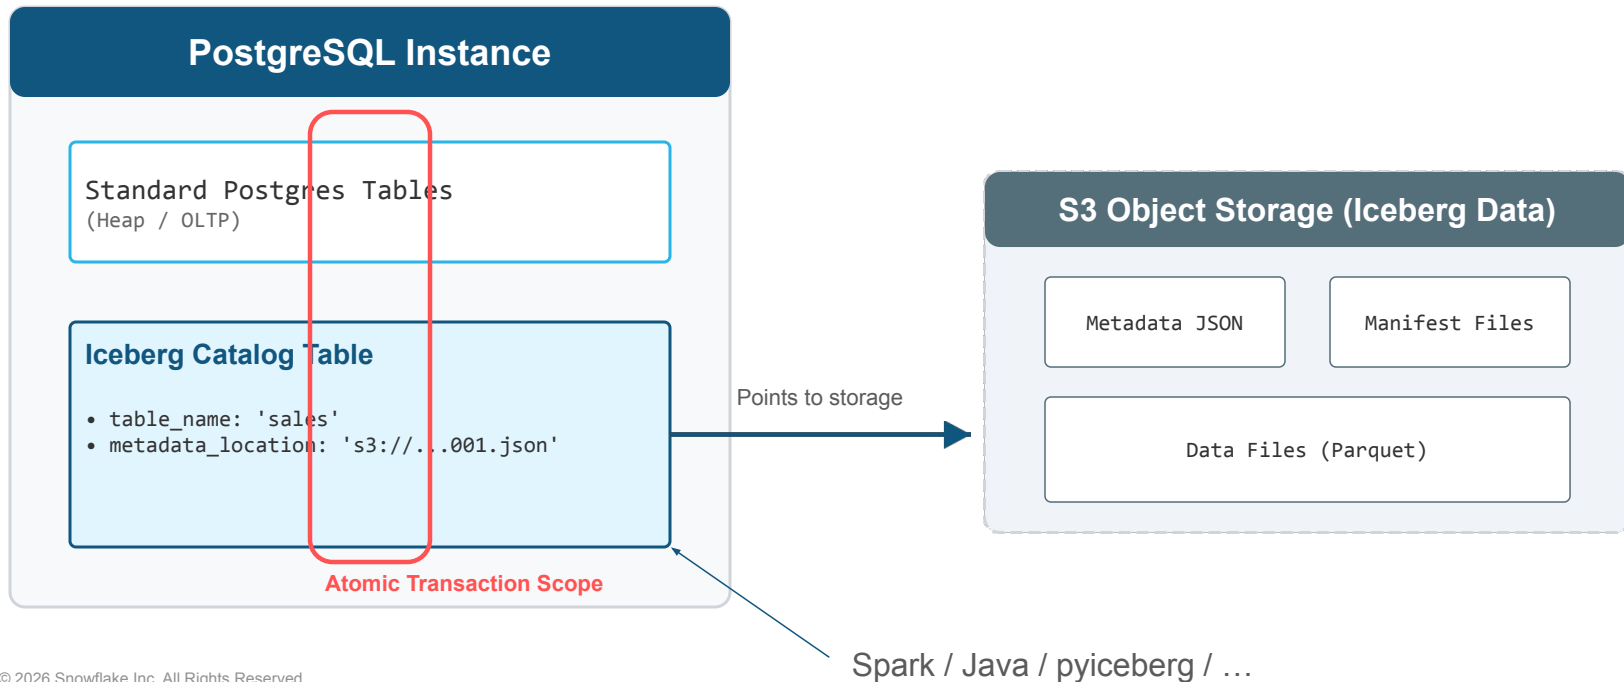


pg_lake architecture



Transactional integrity across table formats

pg_lake ensures atomic commits by managing the Iceberg catalog inside Postgres.



Example: Event queue for Iceberg using pg_cron

Periodically flush a queue into Iceberg with pg_cron.

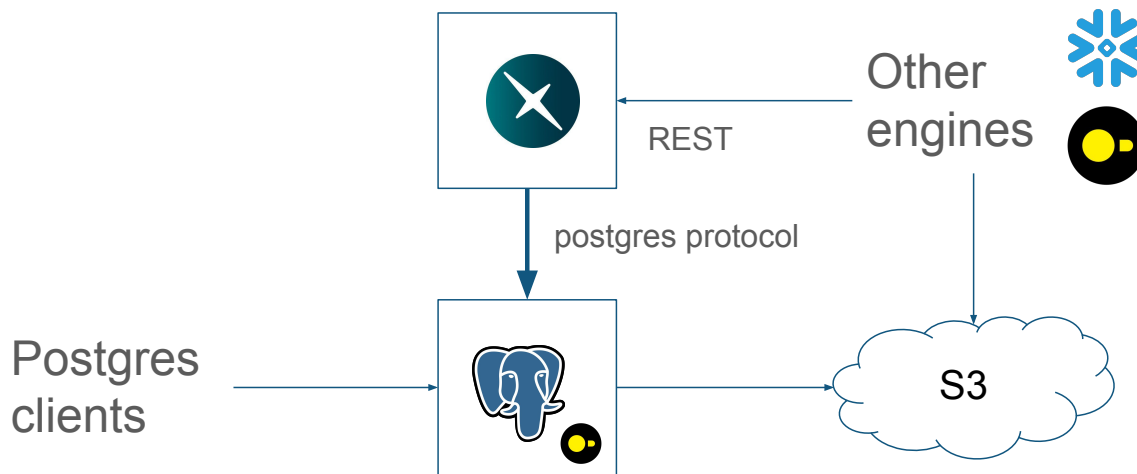
```
create table events_iceberg (  
    event_id bigserial, event_time timestamptz default now(), ...  
)  
using iceberg;  
  
create table events_queue (like events_iceberg);  
  
select cron.schedule('10 seconds', $$  
    with del as (delete from events_queue returning *)  
    insert into events_iceberg select * from del  
    $$);
```



New frontiers: Polaris on pg_lake

Polaris is an Iceberg REST API server that uses Postgres as a metadata store.

Sync to/from pg_lake metadata in the backend in transactions.



github.com/Snowflake-Labs/pg_lake

marco.slot@snowflake.com