



HYDRA

DuckDB-Powered Postgres

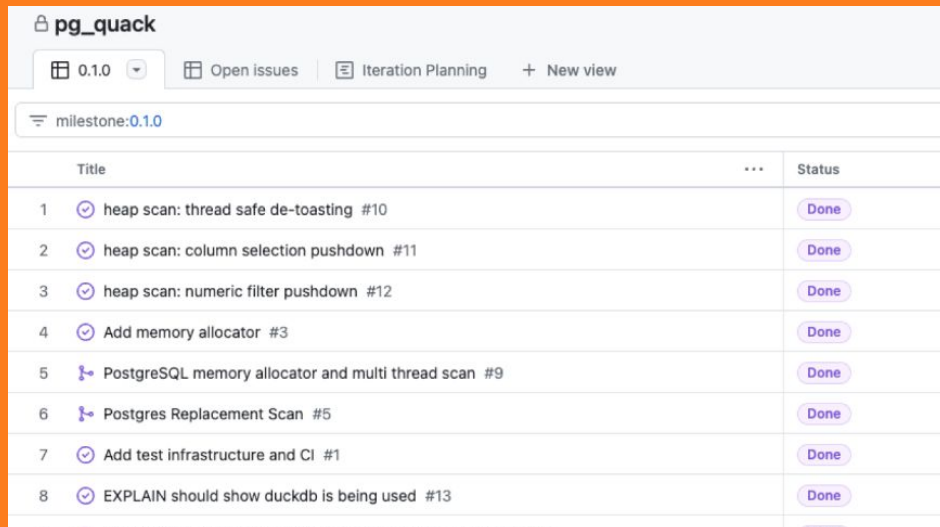
August 2024



What's pg_duckdb?

pg_duckdb is an open source (MIT licensed) Postgres extension that embeds DuckDB's state-of-the-art analytics engine and features within Postgres.

Using **pg_duckdb** is the easiest way to build real-time applications and analytics with Postgres.



The screenshot shows the GitHub project page for **pg_quack**. The page is filtered by the milestone **0.1.0**. It displays a list of 8 issues, all of which are marked as "Done".

Title	Status
1 <input checked="" type="checkbox"/> heap scan: thread safe de-toasting #10	Done
2 <input checked="" type="checkbox"/> heap scan: column selection pushdown #11	Done
3 <input checked="" type="checkbox"/> heap scan: numeric filter pushdown #12	Done
4 <input checked="" type="checkbox"/> Add memory allocator #3	Done
5 <input checked="" type="checkbox"/> PostgreSQL memory allocator and multi thread scan #9	Done
6 <input checked="" type="checkbox"/> Postgres Replacement Scan #5	Done
7 <input checked="" type="checkbox"/> Add test infrastructure and CI #1	Done
8 <input checked="" type="checkbox"/> EXPLAIN should show duckdb is being used #13	Done

What can pg_duckdb do?


- SELECT queries executed by the DuckDB engine can directly read Postgres tables.
- Read parquet and CSV files from object storage (AWS S3, Cloudflare R2, or Google GCS).
- Enable the DuckDB Iceberg extension using `SELECT duckdb.enable_extension('iceberg')` and read Iceberg files with `iceberg_scan`.
- Write a query — or an entire table — to parquet in object storage.
- Read and write to Parquet format in a single query
- Query and JOIN data in object storage with Postgres tables, views, and materialized views.
- Create indexes on Postgres tables to accelerate your DuckDB queries

pg_duckdb repo

duckdb/pg_duckdb



Let's try it out



DEMO

To wrap up

duckdb/pg_duckdb

Website: [Hydra.so](https://hydra.so)

@hydradatabase

@JoeSciarrino



THANK YOU



quack
quack