
Watershed
Get to zero carbon, fast.


InsightGlobal

Walmart

DOORDASH

BlackRock

Figma

BAILLIE GIFFORD

WISE

airbnb

Match
Group

Grover

Sonder
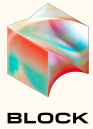
BLOCK

Misfits Market

athenahealth

ZAXBY'S

monzo

CSWR
Central States Water Resources

BEAUTYCOUNTER

EVERLANE

BOMBAS

TaskUs™

BOOM

Klarna.

shopify

data
iku

BOB'S DISCOUNT FURNITURE

WARBY PARKER

Spotify

GOJO
Purell
GOJO, MAKERS OF PURELL™

okta

sweetgreen

Oakley Capital

FORTRESS

General
Mills

stripe

servicenow

sunrun

THOMABRAVO

SKIMS

kainos

BainCapital

Wolt

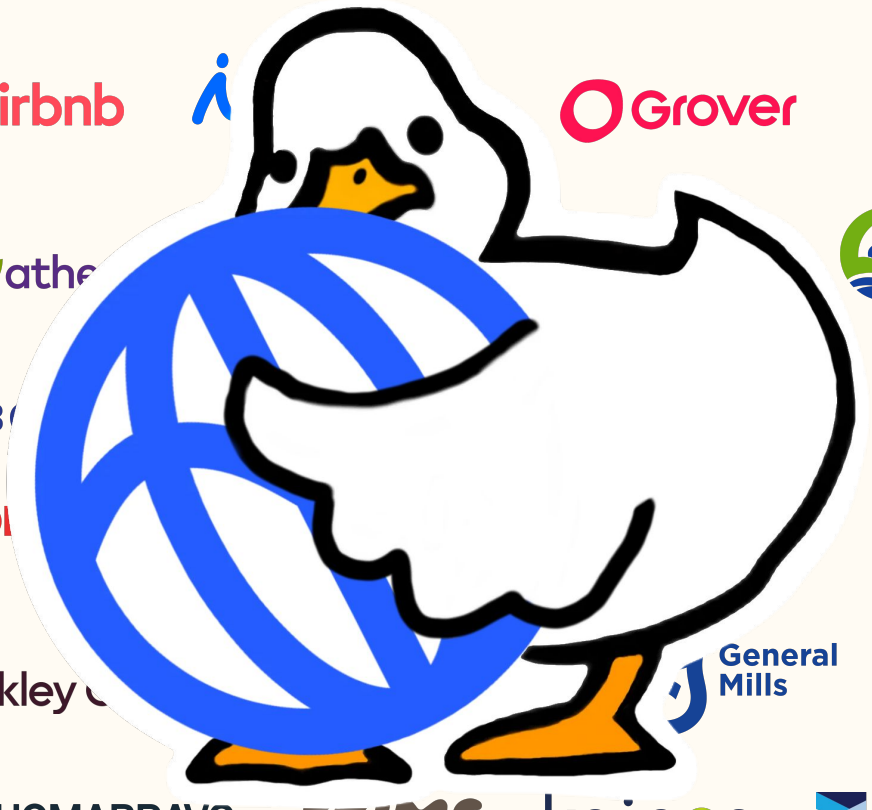
InsightGlobal
Walmart
DOORDASH
BlackRock
Figma
BAILLIE GIFFORD
WISE
airbnb
Grover
Sonder
BLOCK
Misfits Market
CSWR
Central States Water Resources
BEAUTYCOUNTER
EVERLANE
Klarna.
shopify
data iku
GOJO
Purell
GOJO, MAKERS OF PURELL™
okta
sweetgreen
General Mills
stripe
servicenow
sunrun
THOMABRAVO
SKIMS
kainos
BainCapital
Wolt

Watershed

Carbon footprint data

User needs: fast aggregates

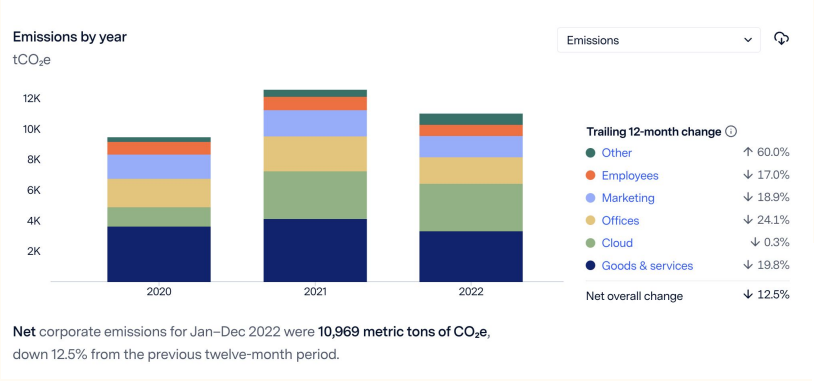

Medium-sized data

- 12% of customers have footprints with > 1m rows
- Largest customer has >15m rows → ~750mb parquet

Medium-sized data

- 12% of customers have footprints with > 1m rows
- Largest customer has >15m rows → ~750mb parquet

Medium-sized data

- 12% of customers have footprints with > 1m rows
- Largest customer has >15m rows → ~750mb parquet
- Fits on one machine, but non-trivial to query performantly

Existing solution: Postgres

Existing solution: Postgres

- 162GiB table / 252GiB of Postgres database size

Existing solution: Postgres

- 162GiB table / 252GiB of Postgres database size
- Painful migrations
- Arbitrary analytic queries don't scale

Architecture


browser
Google Cloud Storage
Carbon footprint parquet files

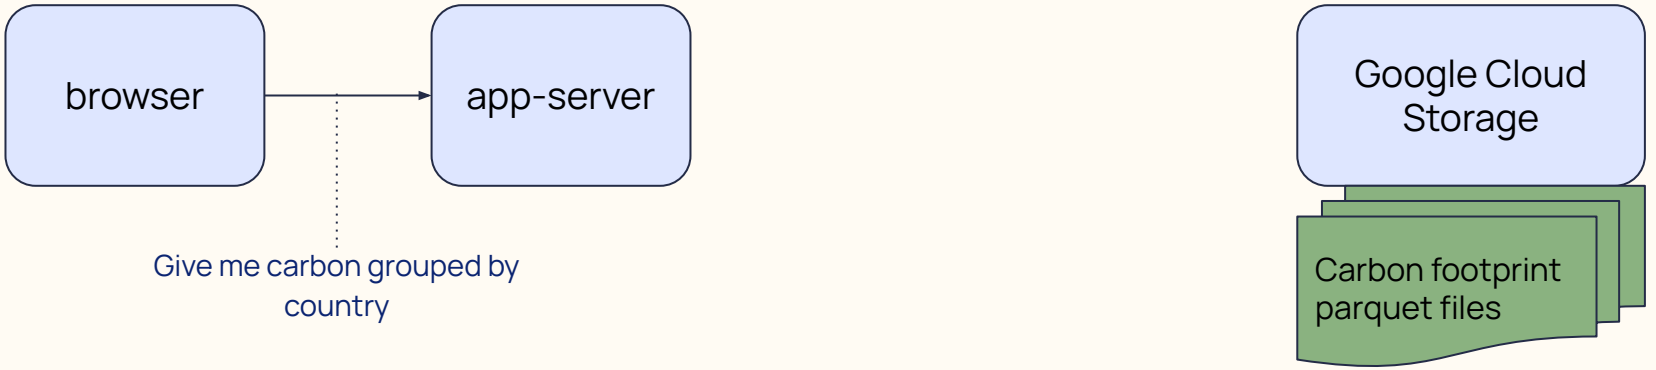
browser
app-server
Give me carbon grouped by country
Google Cloud Storage
Carbon footprint parquet files

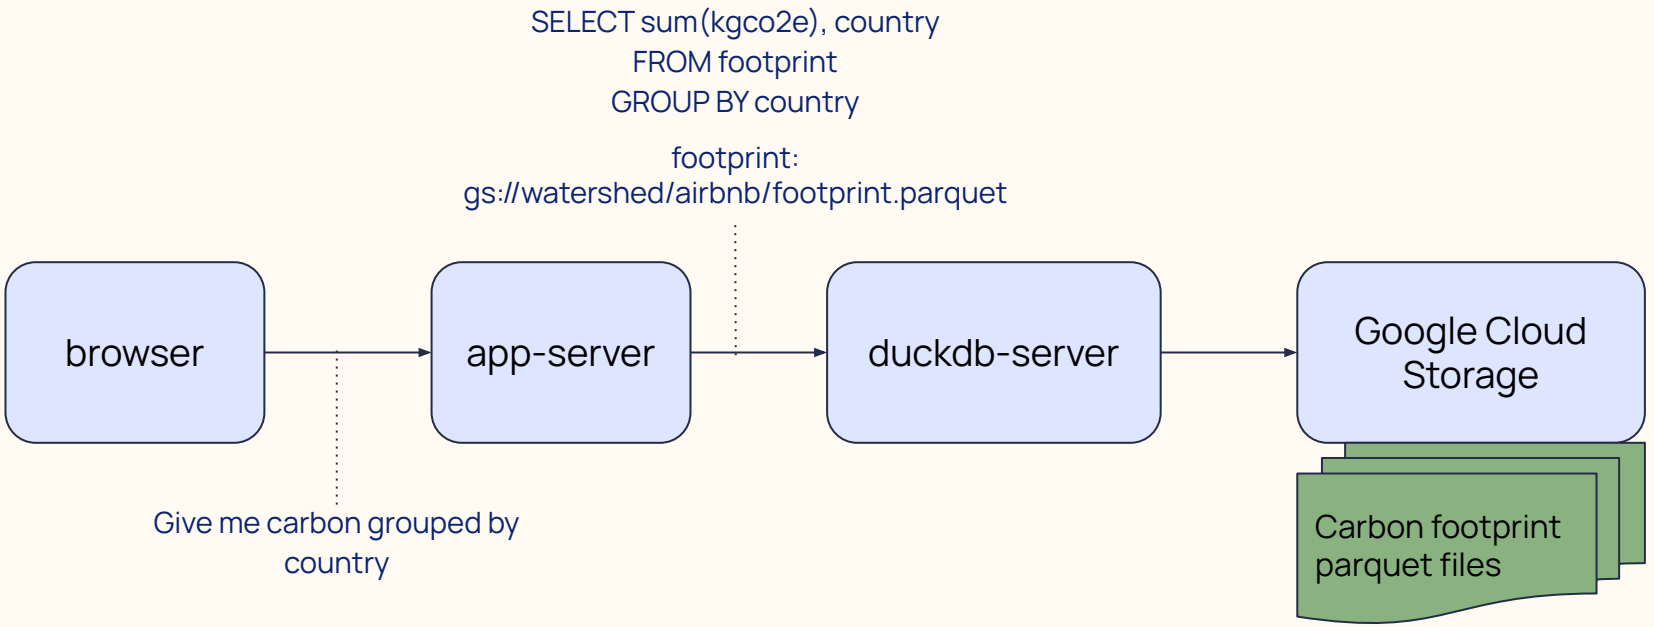
SELECT sum(kgco2e), country
FROM footprint
GROUP BY country
footprint:
gs://watershed/airbnb/footprint.parquet
browser
app-server
duckdb-server
Google Cloud Storage
Carbon footprint parquet files
Give me carbon grouped by country

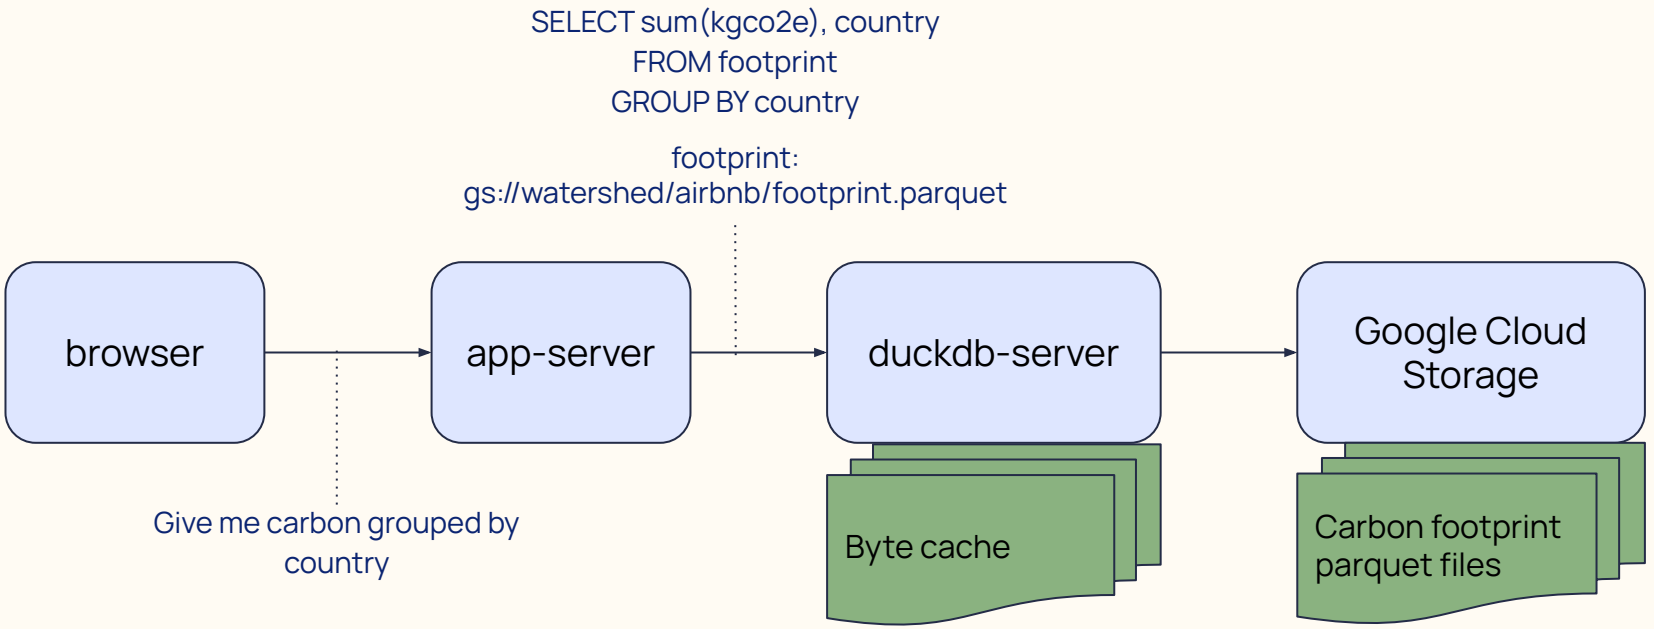
SELECT sum(kgco2e), country
FROM footprint
GROUP BY country
footprint:
gs://watershed/airbnb/footprint.parquet
browser
app-server
duckdb-server
Google Cloud Storage
Byte cache
Carbon footprint parquet files
Give me carbon grouped by country

Why we love this

- It's fast!
  - 10x faster than Postgres with lots of indexes
  - 100s of ms for our P99 data size
- No more giant table!
- Less adhoc caching
- Parquet 💜

Other uses

- Data pipeline: activity data → carbon footprint data
- Query any parquet file
- ~75k duckdb-server queries per day

Other uses

- Data pipeline: activity data → carbon footprint data
- Query any parquet file
- ~75k duckdb-server queries per day

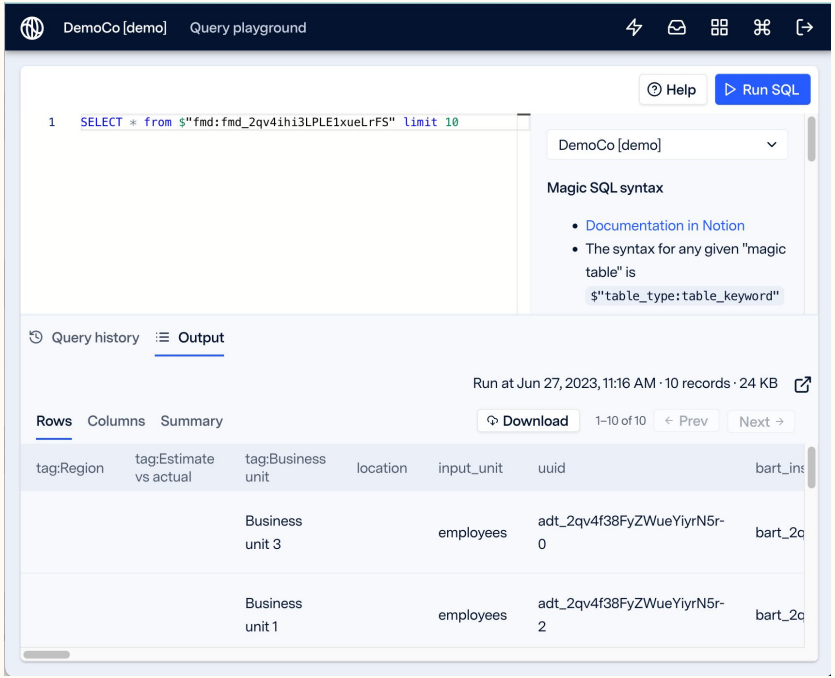

Thank you!

jessica@watershed.com

Watershed