

Tabular Database Systems

⑤

Database-External Data in Parquet Files 🗄️


March 17, 2026

Torsten Grust
Universität Tübingen, Germany

1 | Columnar Compressed Data Storage Outside the DBMS: Parquet

Fragility, space requirements, and parsing effort render CSV files problematic for applications that read/write GBs or TBs of data.

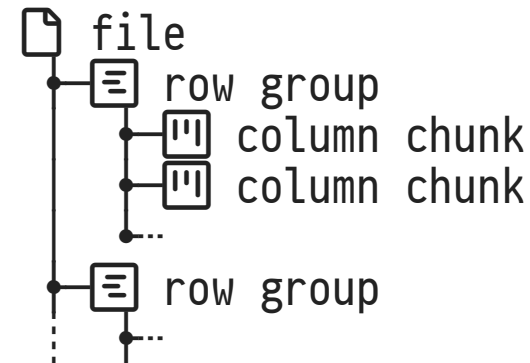

Alternative **database-external file formats** have been developed to address these issues. Among these, **Parquet**¹  is widely used.

- Developed in the open since 2013 (initiated by Twitter *et al*).
- Stores **columns of typed data**.
- Built-in **compression** (on file and column levels).
- Incorporates **rich metadata** that supports projection and selection pushdown.
- Supported by libraries for a wide range of programming languages.
Directly readable/writable by DuckDB .


¹ [Apache Parquet](#)  (sponsored by the Apache Software Foundation): open source, column-oriented data file format designed for efficient data storage and retrieval.

CSV vs. Parquet

CSV	Parquet
monolithic row-oriented (lines separated by \backslash) plain text, uncompressed untyped (requires parsing) no metadata (optional header row)	split into horizontal row groups column-based (chunk-by-chunk) supports compression: <ul style="list-style-type: none"> • file-level (gzip, zstd, ...) • column-level (dictionary, RLE, ...) typed: <ul style="list-style-type: none"> • scalar (INT32/64, FLOAT, BYTE_ARRAY, ...) [• nested records] metadata for file/row groups/columns: <ul style="list-style-type: none"> • file format version • min/max/count statistics • cardinality, (un)compressed byte sizes




A Sketch of Parquet's Storage Format



2 : Pushing Projection and Filtering Down into Parquet Reading

Parquet's file structure + metadata enables readers (like DuckDB's `PARQUET_SCAN` operator) to **only access relevant data subsets**.

- **Projection pushdown:** Exploit **column-based layout**. In each row group, use `data_page_offset` to navigate the file to only read required column chunk(s).
- **Filter pushdown:** Exploit **statistics metadata**. Entirely skip row group if `stats_min/stats_max`² for column `c` indicate that filter predicates like `lo ≤ c ≤ hi` will always fail (///).




² In the DB research literature, the `(stats_min, stats_max)` pairs in all row groups are sometimes collectively referred to as **zone map** (for column `c`).

3 : Optional in Parquet Files: Bloom Filters

If column values are randomly shuffled, all values may occur in all row groups and the effectiveness of zone maps is largely lost.

- **Bloom filters** are *compact* data structures that *over-approximate* the set D of distinct values³ in a row group.
- Given the question $x \in D?$, Bloom filters either respond with
 - *definitely no* (👎) or
 - *probably yes* (👍 _{$p\%$} , where p is a false positive rate).

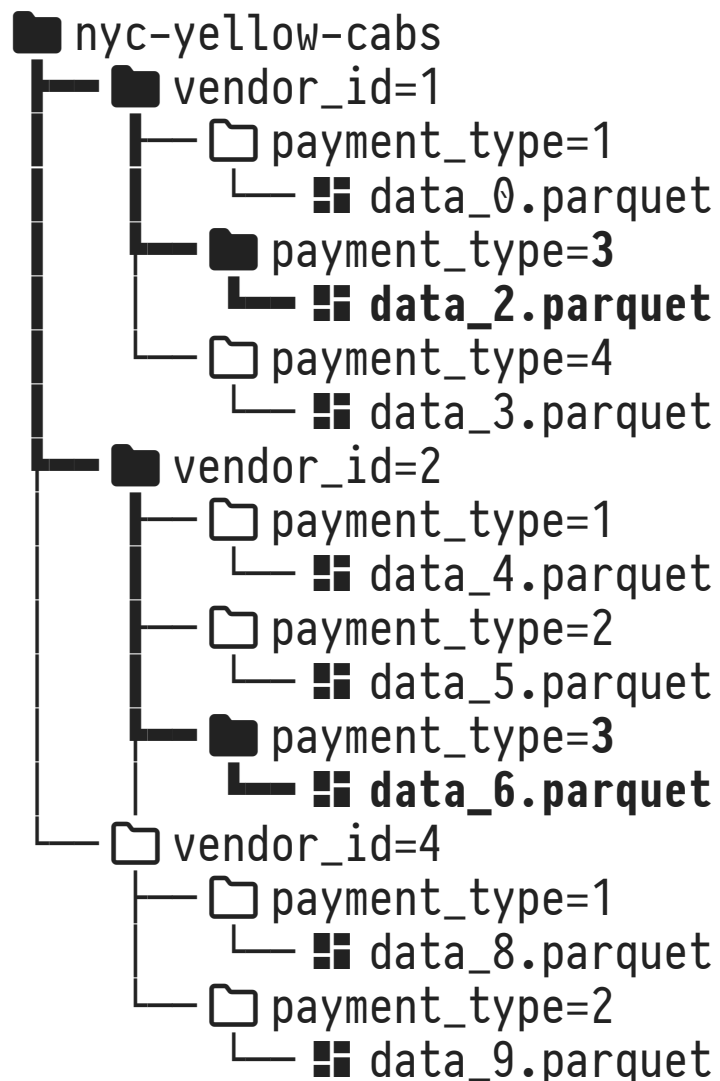


#026

³ Set D is the **active domain** of the row group. Example: if the values in a column have type `int32` with domain $[-2^{31}, 2^{31})$, we typically have $D \subseteq [-2^{31}, 2^{31})$ (read \subseteq as “significantly smaller”).

4 : Partitioned Files and Filter Pushdown (Hive Partitioning)

If data volumes get truly large, it may make sense to **partition rows** into a family of files. Which partition will a row belong to?



- Choose $n \geq 1$ columns c_1, \dots, c_n as partition criteria (left: $n=2$, `vendor_id` + `payment_type`).
- Arrange partitions in **tree hierarchy**:
 - Depth of tree: $n+1$.
 - Width of tree \leq product of the size of active domains of columns c_i .
- Inner nodes: OS directories, leaf nodes: data (Parquet or CSV).
- Filter query: traverse relevant subtrees/read relevant data files only (left: \equiv `payment_type = 3`).
- Particularly useful when I/O is slow.