

PostgreSQL Summary Stats and Window Functions

Aggregate window function: A traditional aggregate (SUM, AVG, COUNT, MIN, MAX) applied as a window function so it returns a per-row aggregate computed over the rows in the window or partition

COALESCE: A scalar function that returns the first non-NULL value from its argument list, commonly used to replace NULLs from ROLLUP/CUBE or LAG/LEAD with a default

CREATE EXTENSION: A PostgreSQL command that enables an extension (such as tablefunc) so its additional functions and features become available in the database

CROSSTAB: A PostgreSQL tablefunc extension function that pivots rows into columns by turning distinct values of a column into output columns, producing a cross-tabulated result

CTE (Common Table Expression): A named temporary result set defined with WITH that you can reference within a single SQL statement to simplify queries and improve readability

CUBE: A GROUP BY modifier that produces all possible combinations of group-level aggregations for the listed columns, yielding every subtotal plus the grand total

DENSE_RANK: A ranking window function that assigns equal ranks to tied values but does not leave gaps, so ranks increase consecutively after ties

AWS Lambda: A serverless compute service that runs code in response to events and automatically manages the underlying compute resources, charging only for actual execution time

Data warehouse: A centralized computer system designed to gather, integrate, store, and make large amounts of organizational data available for reporting and analysis

LAG: A relative fetching window function that returns the value of a column from n rows before the current row, returning NULL if that prior row does not exist

LAST_VALUE: An absolute fetching window function that returns the last value in the window or partition and often requires extending the frame (e.g., RANGE/ROWS BETWEEN UNBOUNDED PRECEDING AND UNBOUNDED FOLLOWING) to get the true last value

LEAD: A relative fetching window function that returns the value of a column from n rows after the current row, returning NULL when no following row exists

Moving average: A time- or order-based window calculation that averages a value over a specified frame of current and preceding (or following) rows to smooth short-term fluctuations

NTILE: A window function that divides ordered rows into n approximately equal buckets (tiles) and assigns each row a bucket number for paging or quantile classification

ORDER BY (within OVER): A subclause of OVER that defines the ordering of rows inside the window and determines how functions like ROW_NUMBER, RANK, and frame-based aggregates are applied

OVER clause: The clause that identifies a function as a window function and defines the window by specifying subclauses such as PARTITION BY, ORDER BY, and frame clauses

PARTITION BY: An OVER subclause that splits rows into separate partitions (groups) so the window function is applied independently to each partition

RANGE BETWEEN: A frame-specification that defines the window in terms of logical ranges of the ORDER BY expression and treats duplicate ORDER BY values as a single range entity

RANK: A ranking window function that assigns equal ranks to tied values and leaves gaps in subsequent ranks (e.g., ties at rank 2 cause the next rank to be 4)

ROLLUP: A GROUP BY modifier that produces hierarchical subtotal rows for specified grouping columns, useful for generating group-level and grand totals in a concise way

ROW_NUMBER: A window function that assigns a unique sequential integer to each row within the window, based strictly on the specified ordering

ROWS BETWEEN: A frame-specification that defines the window in terms of a physical number of rows before or after the current row, e.g., ROWS BETWEEN 2 PRECEDING AND CURRENT ROW

Running total (cumulative sum): A SUM used as a window function over an ordering and typically a frame that starts at the partition beginning and ends at the current row to produce a cumulative sum

STRING_AGG: An aggregate function that concatenates values from multiple rows into a single string using a specified separator, useful for compressing lists into one row

Window function: A SQL function that performs a calculation across a set of table rows related to the current row without collapsing rows, producing a result for each input row