redgate

SQL Server terms translated to PostgreSQL

SQL Server		PostgreSQL
Instance	The running server process that hosts databases, users, and accompanying processes.	Cluster
User	Principle for connecting to the server for administration or running queries. While PostgreSQL does have a CREATE USER command, it is simply an alias to CREATE ROLE WITH LOGIN.	Role w/Login
Row	Tuple is used as a synonym for row in most PostgreSQL documentation and discussions. A Tuple technically refers to the schema of the row. Therefore, a row is an individual instance of a Tuple.	Tuple
Transaction Log	Ongoing transaction storage for A.C.I.D. compliance.	Write Ahead Log
Buffer Pool	Reserved memory that stores frequently accessed data for quick retrieval. In PostgreSQL, this is configured statically at the cluster level and should be adjusted to match your workload and available server memory.	shared_buffers
Query Store	PostgreSQL does not currently include the ability to store query plans and execution metrics over time like SQL Server Query Store. Third-party tools like SQL Monitor can help with this.	NO EQUIVALENT
Query Plan Cache	PostgreSQL does not cache query plans like SQL Server. Instead, plans are compiled with each execution, unless you use PREPARE. However, prepared statements only last for the lifetime of the session.	NO EQUIVALENT
LOB	Like SQL Server, PostgreSQL needs to store large data values that cannot fit onto an 8kB page off table. The method for doing this is called The Oversized-Attribute Storage Technique, or TOAST.	TOAST
[Table One]	Qualifying reserved words, identifiers, and objects.	"Table One"
TOP	Limit the number of rows returned from a query.	LIMIT
GETDATE()	Retrieve the current date and time in server-local time.	NOW()
DATEADD (day, 1, '2022-01-01')	Dates, Timestamps and INTERVAL types can be added and subtracted directly in PostgreSQL, rather than using a special function.	'2022-01-01'::date + '1 day'::interval
Indexed View	On-disk, snapshot representation of a VIEW for fast runtime query. Both servers allow you to create indexes and refresh the snapshot.	Materialized View
Inline code	In SQL Server, all SQL is executed in a T-SQL shell by default, allowing procedural features like variables and loops. PostgreSQL, however, must be told explicitly to interpret a block of SQL as procedural using an anonymous code block with a specified language like pl/pgsql.	Anonymous Code Blocks (DO blocks)
Actual Query Plan	In PostgreSQL there is no official graphical IDE similar to SSMS. Options exist, but few provide a visual query plan explorer. Instead, PostgreSQL returns a text-only plan represented as a hierarchy of nodes (operators). Tools like SQL Monitor provide options for seeing these in a visual way.	EXPLAIN ANALYZE
CROSS APPLY	CROSS APPLY in SQL Server is a very helpful way to manage the execution order of a query and to return data that would otherwise require functions and procedures. In PostgreSQL, a LATERAL JOIN provides similar functionality and usefulness.	INNER JOIN LATERAL
BULK INSERT	Efficiently inserting large amounts of data into a database can be challenging. SQL Server and PostgreSQL have options to make this easier. COPY is typically the fastest tool for inserting formatted data into a PostgreSQL database. Many popular SDKs implement the COPY binary protocol for efficient processing.	COPY

Want to know more about PostgreSQL basics?

Join PostgreSQL expert, Ryan Booz, for PostgreSQL 101 - the webinar series that introduces the what, how & why of PostgreSQL.

Learn more