

Pipelined table functions

A feature introduced in Oracle 9i in PL/SQL is the pipelining of table functions, allowing result sets to be passed from one function to another, letting us set up a sequence of transformations with no table to hold intermediate results. Rows of the result set can be returned a few at a time, reducing the memory overhead for producing large result sets within a function.

Table functions are functions that produce a collection of rows (either a nested table or a varray) that can be queried like a physical database table. You use a table function like the name of a database table, in the FROM clause of a query. A table function can take a collection of rows as input. An input collection parameter can be either a collection type (such as a VARRAY or a PL/SQL table) or a REF CURSOR. Data is said to be pipelined if it is consumed by a consumer (transformation) as soon as the producer (transformation) produces it, without being staged in tables or a cache before being input to the next transformation. Pipelining enables a table function to return rows faster and can reduce the memory required in caching a table function's results. A pipelined table function can return the table function's result collection in subsets. The returned collection behaves like a stream that can be fetched from on demand. This makes it possible to use a table function like a virtual table.