

## External Tables: Frequently Asked Questions

### 1. What are External Tables?

External tables let you read and query data that resides in flat files, as though the files were in an ordinary Oracle table. It's more convenient to use External Tables instead of SQL\*Loader to load data from flat files into the database because you don't need to create a control file or a staging table to upload and then transform the data. It's all done in the same step. This lets the user concentrate on the queries to perform the transformation, thus saving time during the load process.

To use External Tables, you run: `INSERT INTO transaction_table SELECT * FROM external_table` to copy the data into the transaction table and do any transformations that are necessary.

### 2. How does External Tables work?

Oracle uses both SQL\*Loader and Data Pump access drivers for External Tables. The default access driver is `ORACLE_LOADER`, which allows the reading of data from external files using the Oracle loader technology. By providing the database with metadata describing an external table, the database is able to expose the data in the external table as if it were data residing in a regular database table. The external data can be queried directly and in parallel using SQL.

The second access driver, `ORACLE_DATAPUMP`, lets you unload data—that is, read data from the database and insert it into an external table, represented by one or more external files—and then reload it into an Oracle Database. Similar to the traditional export and import utilities, the file is in Oracle's proprietary format and is not human-readable.

This feature first appeared back in Oracle 9i Release 1 as an alternative to SQL\*Loader. It provided the ability to read and import flat file data without first loading it into a staging table. In the 9i release, it was only possible to create read-only external tables. Writeable external tables were introduced in Oracle Database10g and made it possible to unload data from the database while at the same time creating an external table to read the data back again.

### 3. What can you load the data from?

Disk or named pipe.

### 4. When should you use External Tables versus SQL\*Loader for best load performance?

The record parsing of External Tables and SQL\*Loader is very similar, so normally there is not a major performance difference in the same record format. However, External Tables may be more appropriate in the following situations:

- You want to transform the data as it is being loaded into the database. (You have the power of SQL versus the more limited SQL\*Loader capabilities to do complex transformations.)
- You want to load very large files with large quantities of data. It's very convenient to use transparent parallel processing to do massively parallel loads (100 parallel slaves or more).

Note that you cannot use External Tables to load a file over the network. The file to be loaded must be on the server itself.

5. Are there any new enhancements to External Tables in Oracle Database 11g?

Yes, there is now support for compressed data (read and write), encrypted data (read and write), and the ability to read from a preprocessed input data stream (available from 11.1.0.7). External Tables can be preprocessed by user-supplied preprocessor programs. By using a preprocessing program, users can use data from a file that is not in a format supported by the driver. For example, a user may want to access data stored in a compressed format. Specifying a decompression program for the ORACLE\_LOADER access driver allows the data to be decompressed as the access driver processes the data.