

SQL*Loader: Frequently Asked Questions

1. What is SQL*Loader?

SQL*Loader (sqlldr) is a bulk loader utility used for moving data from flat files into Oracle database tables. It supports various load formats and multi-table loads.

2. What is the SQL*Loader control file?

The control file is a text file that contains DDL instructions. It tells SQL*Loader where to find the data, how to parse and interpret the data, where to insert the data, etc. It is not the same as a database control file.

3. Can you selectively load and modify the data as the database gets loaded?

Yes, you can selectively load data, specify the data character set and manipulate the data with SQL functions.

4. Can you load data from multiple data files into multiple tables at the same time?

Yes.

5. What can you load the data from?

Disk or named pipe.

6. When should you use SQL*Loader versus External Tables 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 load data remotely. You use SQL*Loader to load a file from a client machine to the server. (External Tables cannot be loaded over the network and only works on the server.)
- Transformations are not required on the data, and the data does not need to be loaded in parallel.
- You want to load data, and additional indexing of the staging table is required.

7. How can you improve the performance of SQL*Loader?

- Do not have any indexes and/or constraints (primary keys) on your load tables during the load process.

- Add the following option in the command line: `DIRECT=TRUE`. This will bypass most of the RDBMS processing by using the direct path loader instead of the conventional path loader. However, there are some cases when you can't use direct load. These restrictions can be obtained from the Oracle Server Utilities Guide.
- Use fixed width data rather than delimited data. For delimited data, each record needs to be scanned for the delimiter.
- Try to avoid character set conversions as conversions are both time and cpu intensive.
- For conventional path, use the `READSIZE` and `BINDSIZE` parameters. `READSIZE` will grab larger chunks of data per read system call. The `BINDSIZE` parameter specifies the size of the bind array, which in turn specifies the number of rows which will be loaded per batch.

8. Is there a SQL*Unloader to download data to a flat file?

No, Oracle does not supply any data unload utilities.

9. Where can I find more information on SQL*Loader?

Refer to Oracle documentation or the following O'Reilly book: "Oracle SQL*Loader: The Definitive Guide" by Jonathan Gennick and Sanjay Mishra (ISBN 1-56592-948-9).