

## Performance: "pl/sql" vs. "OS b-tree files"

**Source:** <http://coding.derkeiler.com/Archive/Java/comp.lang.java.databases/2004-03/0252.html>

---

**From:** José Antonio Cuesta ([joseantonio.cuesta\\_at\\_ono.es](mailto:joseantonio.cuesta_at_ono.es))

**Date:** 03/29/04

Date: Mon, 29 Mar 2004 15:31:57 GMT

Hi,

we're developing an application in which performance is a very important issue.

Database is Oracle 9.2, application language is Java 1.4.2., OS is Solaris 8 (RAM 8GB, 4x750Mhz CPU).

The process needs to query for every object several times (around 10) the same configuration tables, which are stored in Oracle. Every process has millions of objects, so the number of consults made to Oracle would be rather high.

No INSERT statements would be made, only SELECTs.

The two alternatives are:

- \* wrap the SQL statements in pl/sql packages to minimize the number of SELECT statements made to Oracle. The packages would be called using JDBC. The config tables would be set as CACHED, and perhaps index-organized.

- \* dump all the configuration tables in a OS file and, using Java classes based on B-Trees, build a B-Tree file and make all the searches using Java methods against this B-tree file. This file would be cached in memory by the OS.

It is supposed that the Java methods querying in-memory data would beat the pl/sql procedures querying the Oracle tables... Does it worth to dump all the information from Oracle to a B-Tree file, having in mind that Oracle uses the same technique to build indexes? Which solution would be more efficient?

Thank you very much for your comments. Any hint would be welcomed.