

Performance Tuning with reference to Oracle

Hardeep Singh

Digital Wealth <http://blog.Hardeep.name>

Copyright: Hardeep Singh, 2009. All rights reserved.

Better understanding of this tutorial can be gained if my basic video tutorial on the same topic is viewed prior to reading this. The URL:

<http://blog.hardeep.name/computer/20090217/sql-tuning/>

It is my basic contention that people need to get off to a certain level of understanding before they can read and work through any document on tuning.

Performance tuning activity can be applied to:

1. An application
2. Process tuning
3. A SQL query

Whenever an application requires tuning, the first step is to go look at the code. If the code is in-optimal, no amount of Oracle tuning will speed up the application. Once that step is taken care of, we can come to Oracle level tuning.

Let's start with **SQL query tuning**.

A **filter** is something that restricts certain rows from appearing in the output. For example

`where employee.emp_type='P'`

However, the following is not a filter, it's a **join**:

`where a.emp_type = b.emp_type`

Some of the common query pitfalls are:

- Using NOT EQUAL TO prevents the use of indexes.
`a.emp_type <> 'P'`
Replace with EQUAL TO, LESS THAN or GREATER THAN wherever possible
`a.emp_type = 'T'`
- Large IN lists also prevent use of indexes
- Functions (eg substring) prevent the use of indexes
- UNION operation involves an internal DISTINCT operation to eliminate all the duplicate rows, and as such is expensive. If the distinctness is not that important for your query, use UNION ALL instead of UNION

If there is a large IN list that cannot be avoided, it may be possible to create a small table, just to hold the values in the IN list. Thereafter, the IN list will become an EXISTS operation which is much faster.

Book-Author example

Book	Author
<ul style="list-style-type: none">- Book ID- Name- Reference Book (Y/N)- Author ID- Publisher	<ul style="list-style-type: none">- Author ID- Author Name
Ordered by Book ID	Ordered by Author ID
10 Pages	3 Pages

Consider the two tables above: the book and the author tables. If the information had to be printed the book table would print on ten pages, and the author table on three.

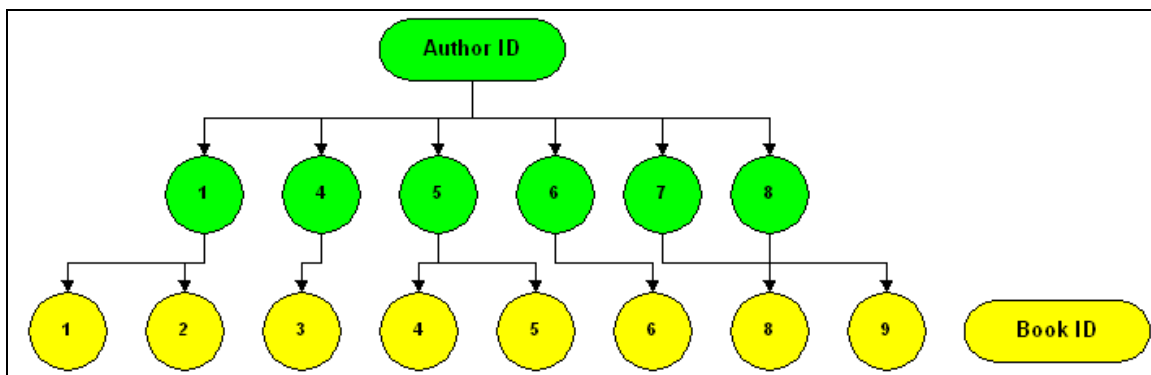
In order to tune queries involving these tables, it's important to imagine how you would search for the information if you were given the table contents printed on paper.

Consider the query:

“I need all the author names who have authored reference books”

We can either start at the book table, or at the author table. If we start with the author table, we can lookup the first author ID, and then look through the entire book table if he has any reference books. The other way is to start with the book table, and when we find a reference book we can check the author name against that. Which one is faster? The second – starting with the book table? Why – because the author table is order by Author ID.

Now let's have a look at how Oracle interprets indexes.



If we had an index on the Book table, on the fields Author ID and Book ID, this is how it would logically look. It can be used if we know both the Book ID and the

Author ID (have both the fields in the where clause of the query). It can also be used if we know just the Author ID, but will be slightly slower (see RANGE SCAN below). However, it cannot be used if we know just the Book ID.

Before executing a query, Oracle creates a **plan** - the order in which it will search the tables, the indexes it will use, the way it will use the indexes etc. Step 1 in tuning is to check the plan for the SQL in question:

<pre> SELECT * FROM ps_voucher a, ps_vchr_acctg_line b WHERE a.voucher_id=b.voucher_id AND a.business_unit=b.business_unit AND a.due_DT='10-JAN-2007' </pre>				
SELECT STATEMENT	Optimizer Mode=CHOOSE	213 K	77605	
TABLE ACCESS BY INDEX ROWID	SYSADM.PS_VCHR_ACCTG_LINE	5	2 K	2
NESTED LOOPS		213 K	172 M	77605
TABLE ACCESS FULL	SYSADM.PS_VOUCHER	43 K	17 M	58093
INDEX RANGE SCAN	SYSADM.PSDVCHR_ACCTG_LINE	1		3

In order to generate the plan, on SQL Plus use:

```
set autotrace on;
```

or if we want to just see the trace without executing the query use:

```
set autotrace traceonly explain;
```

With Toad, you can press **Control – E** to see the plan.

Now looking at the explain plan shown above, what does it mean? The point to note is that explain plan is actually executed from more indented to less indented (inside to outside) and from bottom to top.

Table access full means that the complete table needs to be searched in order to find the information. **Index range scan** means that index will be used, however the filters provided in the WHERE clause may not be 'complete'.

Index full scan means the entire index will be searched.

Table access by index rowid means the rows are being located using an index.

Hash, **Nested loops** and **Merge** are all different types of joins. These are the most commonly seen operations in an explain plan.

More information on the plan: <http://www.orafaq.com/node/1420>

Now that we know how Oracle is executing our query, the next step is to influence the way Oracle runs the query. How to do that?

- ❖ By creating indexes and optimizing the join order of tables
- ❖ By creating **statistics**
- ❖ By adding **hints** to the query

I have already discussed indexes. An example of how to create an index:

```
CREATE INDEX book_author_index ON book (book_id,  
author_id);
```

Statistics refers to the information that Oracle can store about tables and fields on your request that will help it to make more informed decisions regarding the plan. For example, let us say that purchase order table P has a field STATUS which can be 'O' (Open) or 'C' (Close). If you give Oracle a query with this field in the WHERE clause – it will assume that O & C values are equally distributed: 50% each. In truth this may not be the case: O values maybe present only in 10% of the rows. Such information being calculated offline, stored and used during plan making is called statistics.

To create statistics use something like:

```
EXEC DBMS_STATS.gather_table_stats('SYSADM', 'BOOK');
```

More information on statistics: http://www.oradev.com/create_statistics.jsp

A **hint** is written as part of the query, but doesn't affect the output. It only affects the plan.

For example:

```
SELECT /*+ INDEX(A IDX_PURCH_ORD) */ *  
FROM PURCH_ORD A  
WHERE STATUS='O'
```

The hint is written within comments, but with a plus sign to indicate that it's a hint, not a comment. This hint tells Oracle to use the index IDX_PURCH_ORD to lookup the table PURCH_ORD. Most of the time Oracle automatically uses the correct index, however in certain 'tough' situations it needs help.

One of the most useful hints that I have found is the LEADING hint. This hint tells Oracle to start with a particular table from the tables used in the query.

For example:

```
SELECT /*+ LEADING(i) */ po_id, status
FROM PURCH_ORD p, INVOICE i
WHERE invoice='123456' AND
i.po_id=p.po_id
```

In this query, the LEADING hint tells Oracle to start with the INVOICE table, find the PO_ID for the invoice, and use that to lookup the STATUS in the p table.

A good list of all indexes is available here: <http://www.adp-gmbh.ch/ora/sql/hints/index.html>

Now coming to **process tuning**. We already know how to tune an SQL query. The problem that remains now is to find out which query is doing badly, from among all the queries being executed by the process. If the process is small / simple – a simple inspection of the code will give you an idea of the problem SQL. In a slightly larger SQL you may be able to add some code to display the current time before and after each query. However, if that is not possible, or the SQL is very large – we need some automated tracing to find the SQL. One good way to do this is through the **10046 trace**.

To perform this trace the process should execute the following SQL statement as the first step:

```
ALTER SESSION SET EVENTS '10046 trace name context forever,
level 12'
```

Once the process completes, the DBA should be able to check the trace output, and run it through **tkprof** to give you an HTML report. This report gives you a list of the worst queries, the corresponding plans and the waiting times for those queries. It's a wealth of information to pinpoint the problem SQLs and start tuning them.

If the DBA gives you a trace file (from the udump folder on the database server) and not a report, you can use the trace file to generate the report yourself as below:

```
tkprof <trace_file> <output_file> sort=fchela
```

Lets move on now to **application tuning**. This is similar to process tuning except that we need to find the worst performers in the entire application, not just a given process. One way to do that is through the STATPACK analysis. The DBA can help in generating a STATPACK report for a period (say one day) and would give you a report of the worst SQLs running during the period. More information on STATPACK is here: http://download-uk.oracle.com/docs/cd/B10501_01/server.920/a96533/statspac.htm

A quick cheatsheet:

When you get a query to tune, what are the things we can quickly look at to get going?
Here is a quick cheatsheet:

- ❖ Try to limit the output. For example, if we want a list of open orders, we may not be interested in orders from more than 2 years back. Hence, if the `order_date` is an indexed column, we can add:
`order_date >= SYSDATE-700`
- ❖ Have a look at the indexes being used. The index usage may not be optimal.
- ❖ Have a look at the join order of the tables. As a rule, Oracle finds it hard to come up with a good join order where the number of tables is 4 or more. Use the `LEADING` hint if needed, or the `USE_NL` hint.
- ❖ If the query uses multiple business units such as:
`business_unit IN ('NL100', 'NL200', 'NL300')`
it may be faster to run the query separately, once for each business unit. If it must run as a single query, explore connecting the three separate queries through `UNION ALL`
- ❖ If its being run for multiple `FISCAL_YEARS`, the same applies as business unit above
- ❖ If we are joining the tables (or self join), it is important to use all the keys in the join. For example if a table `t1` has three keys:
`t1.key1=t2.key1 AND`
`t1.key2=t2.key2`
This may solve your purpose for the moment (the two keys may be unique in themselves given the scenario), but it will be inefficient.

Additional suggestions:

- **Start out with low hanging fruit**
- **Go after the FULL SCANS**

If you have a system that doesn't perform very well, and the previous step has already been completed, then go after the FULL SCANS. Pick up queries that do FULL SCAN one by one, and tune them. **The timing different may not be much, even after removing the FULL SCAN, but the overall impact on the system will be huge, due to the much reduced IO operations.** The disk is the slowest part of the chain, remember. To check the FULL SCANS going on now, run the SQL

```
SELECT target, a.sid, time_remaining, b.username, b.program, opname
FROM v$session_longops a, active b
WHERE time_remaining>120 AND a.sid=b.sid(+)
```

(you may need to request access to these tables)

- **Reschedule activity**

This is a workaround in cases where performance of the query itself is difficult to improve. If a job with a bad impact on performance can wait till the weekend, better go that route. If it can run during a time when the system is underutilized go for it. If two IO intensive jobs can run sequentially rather than parallelly go for it.

On those same lines, if there are two jobs that run almost simultaneously, it may be prudent to move one of them five minutes above or below the other.

- **Deadlocks**

Deadlocks are the last bane of performance tuning in a production system. Oracle takes some time in detecting deadlocks, and killing the offending process. During this timeframe your database queue can spiral.

These are quite easy to fix: the offending queries are present in the database log. Hence, pick up the queries, search the processes to find where these are running and redesign suitably. Deadlock will go away if the processes **do not try to process the same set of rows at the same time.**

Read more on deadlocks on my blog at <http://blog.hardeep.name/computer/programming/20090622/oracle-deadlocks/>