Why is my Oracle10g Database SLOWER than my Oracle9i Database?

This paper illustrates performance issues I've discovered with some of the features of the Oracle10g database..

This material contains excerpts from my 2-day Advanced Tuning course.

Index Monitoring

Extra indexes are not only used but can significantly slow down batch processing. Oracle will create before images and journaling entries for each update or delete to a table. Oracle also does this for each index affected by the DML operation.

Oracle9i allows for index usage to be easily monitored. Simply issue 'alter index <index name> MONITORING USAGE' (NOMONITORING USAGE turns it back off) to start the monitoring process. Monitor V$_OBJECT_USAGE, in particular, the USED column will be set to 'YES' when the index appears in an explain plan.

View: SYS.V$OBJECT_USAGE

Created: 5/12/2002 4:23:35 PM Updated: 5/12/2002 4:23:35 PM

Columns | Script | Data | Grants | Deps (Uses) | Deps (Used by) | Triggers

Include Updatable?

Column	Data Type	Null?	Updatable	Comments
INDEX_NAME	VARCHAR2 (30)	N		Name of the index
TABLE_NAME	VARCHAR2 (30)	N		Name of the table upon which the ind...
MONITORING	VARCHAR2 (3)	Y		Whether the monitoring feature is on
USED	VARCHAR2 (3)	Y		Whether the index has been accessed
START_MONITORING	VARCHAR2 (19)	Y		When the monitoring feature is turned...
END_MONITORING	VARCHAR2 (19)	Y		When the monitoring feature is turned...

The spool the output from this script: 'select 'alter index ' || index_name || ' monitoring usage; from user_indexes' can be used to trip this monitoring on.

Run a script to check for index usage and drop the unused indexes! Again, a script can be used to drop the unused indexes. Ask Dan for this script.

Things I've Heard…

Index Clustering Factor

Index Clustering factor is the relationship between the Index leaf block and how many different data blocks each leaf block points to. The CBO uses this ratio to determine if it is more efficient to do a full table scan versus a range scan when being asked to perform a range scan.

The example in the slide shows two tables (with permission from Jonathan Lewis). TEST_TABLE1 and TEST_TABLE2 have the same attributes. The data is in significantly different order.

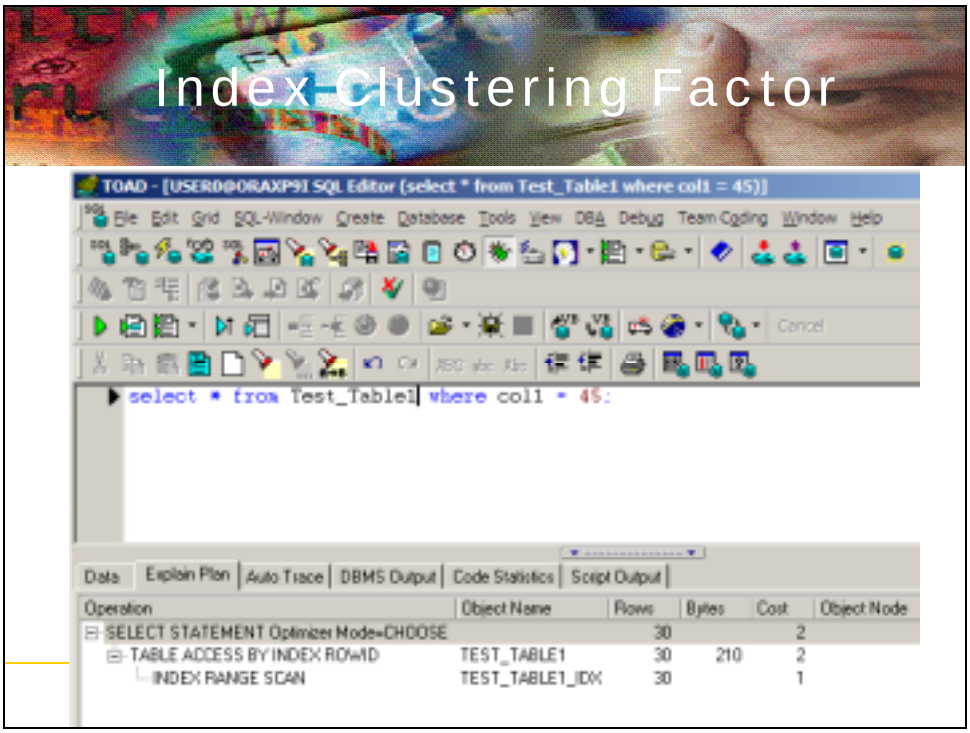

Test Table 1 performs a nice range scan.

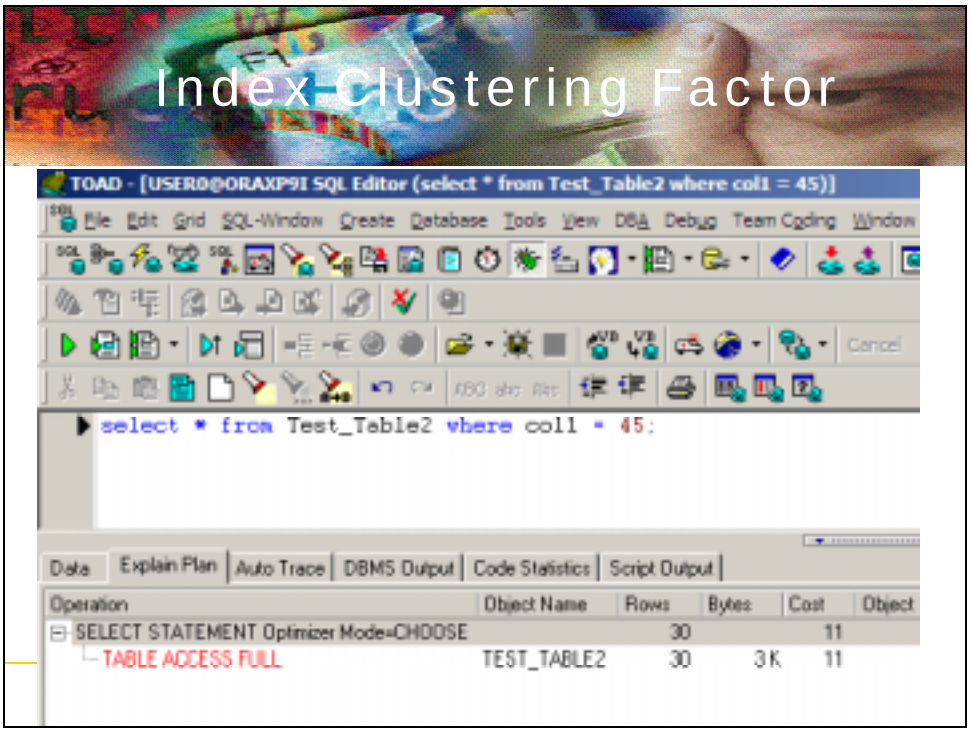

Test Table 2 performs a full table scan at a much higher cost.

Test Table 2 can do a range scan with a hint but at 3 times the cost.

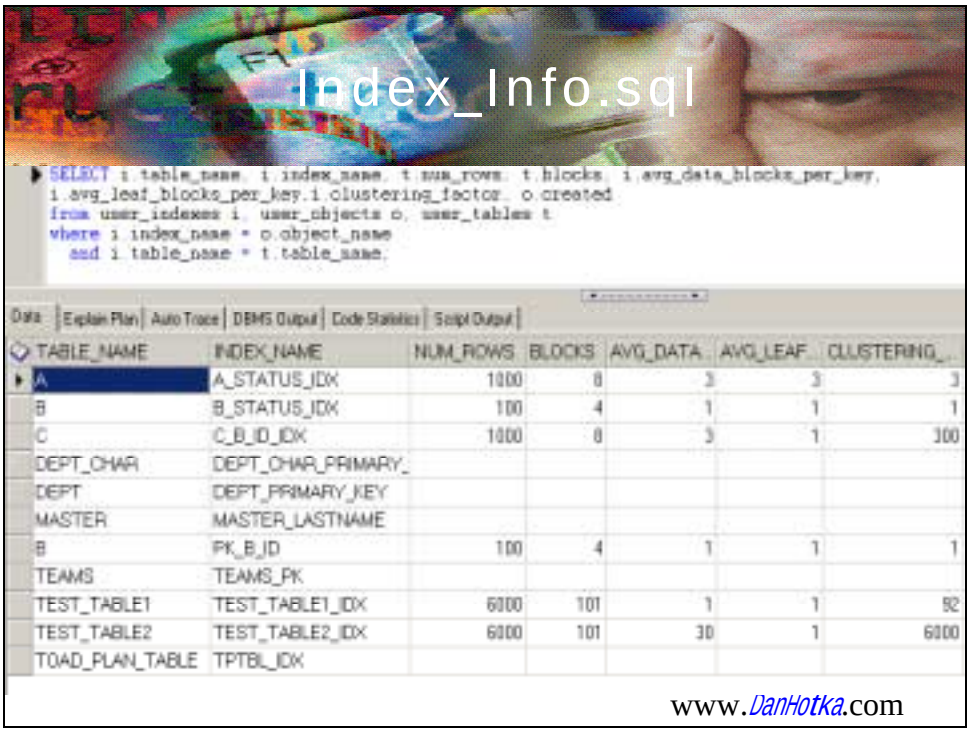

INDEX_INFO shows what is going on. Notice the clustering factor of Test Table 2 is very high. The clustering Factor of Test Table1 is low.

The bottom line is the more the data is in the order of the index key, the more Oracle will use it on range scans.

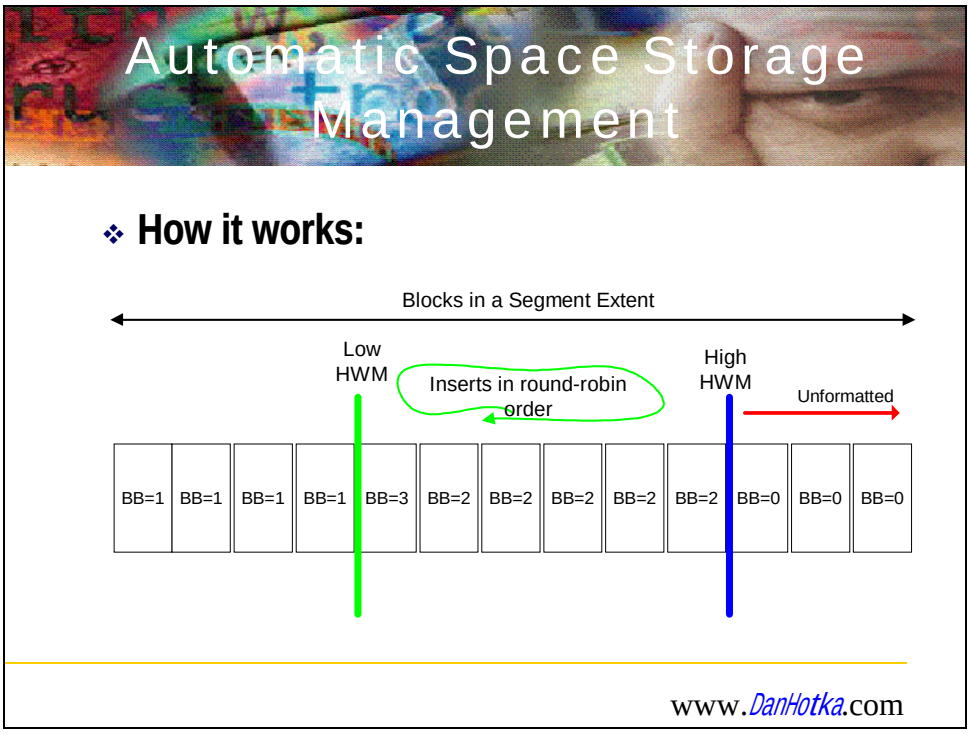

Automatic Space Management is a new tablespace feature for Oracle10g. If used, Oracle does not format the whole data file but just 6 blocks at a time. Oracle will then round-robin insert into these blocks when doing inserts. Doing an insert into a new block will always give a higher clustering factor than just using a regular tablespace.

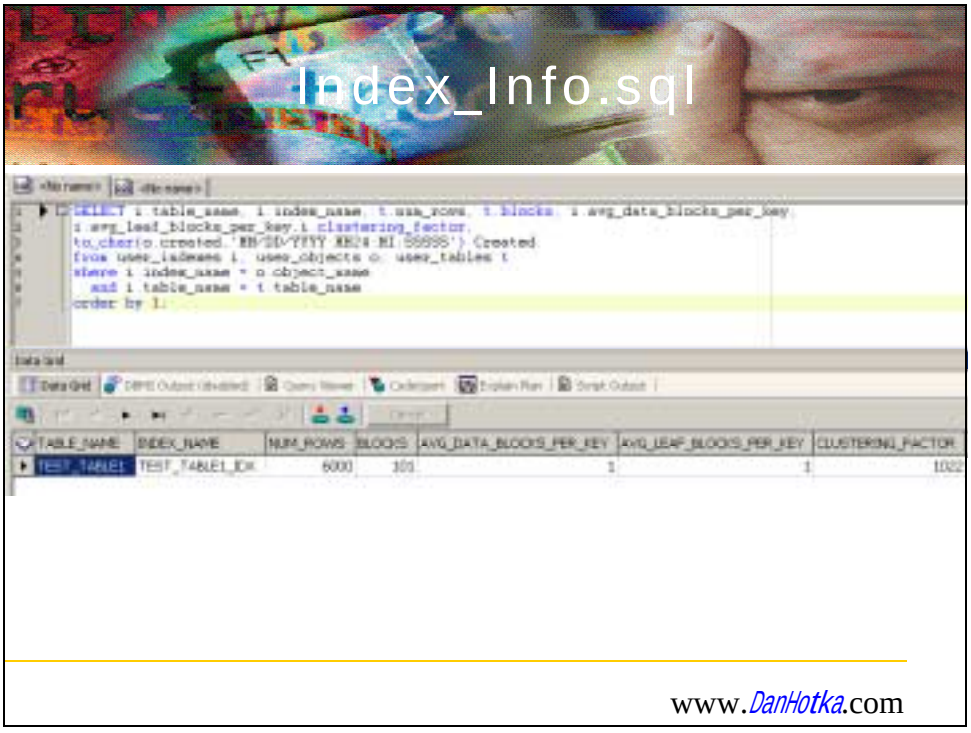

This example shows Test Table 1 against an Oracle10g tablespace using ASM. Notice the higher clustering factor.

ASM is good for RAC and applications with lots of insert activity but possibly at the cost of using indexes on range scan.

Case Study

Extra indexes are not only used but can significantly slow down batch processing. Oracle will create before

Summary

This article fully describes several things to watch and monitor with your Oracle10g database. Clustering factor and the Automated Segment Management feature could be an issue. Too many indexes can adversely affect any Oracle database. Compatibility might help. The case study illustrates adding a lot more memory to the Oracle10g database than what the Oracle9i database required.

Dan Hotka is a Training Specialist who has over 29 years in the computer industry and over 24 years of experience with Oracle products. He is an internationally recognized Oracle expert with Oracle experience dating back to the Oracle V4.0 days. Dan's latest book is the SQL Developer Handbook by Oracle Press. He is also the author of Oracle9i Development By Example and Oracle8i from Scratch by Que and has co-authored 7 other popular books including the Database Oracle10g Linux Administration by Oracle Press. He is frequently published in Oracle trade journals, and regularly speaks at Oracle conferences and user groups around the world. Visit his website at www.DanHotka.com. Dan can be reached at dhotka@earthlink.net .

Dan Hotka - Author/Instructor/Expert
www.DanHotka.com
DHotka@Earthlink.net
515 279-3361