



Official Publication of the Northern California Oracle Users Group

NoCOUG

J O U R N A L

VOL. 18, No. 4 · NOVEMBER 2004

\$15

Take a Tour of Computer History at NoCOUG's Fall Conference



COMPUTER HISTORY MUSEUM

Renew Your NoCOUG Membership for 2005

Membership runs on the calendar year. That means that it's time to renew your NoCOUG membership for 2005. Just go to www.nocoug.org/join.html today!

Grids Are Not Just for Data...

and Other Insights, featuring Craig Shallahamer. You'll also find a review of his popular class on Oracle performance.

Oracle 10g Upgrade Experiences

Learn what happened when a local DBA took the plunge.

Oracle Performance Troubleshooting

Check out this book review by Brian Hitchcock.

Much More Inside . . .

Oracle Performance Troubleshooting

Book reviewed by **Brian Hitchcock**



Oracle Performance Troubleshooting with Dictionary Internals SQL & Tuning Scripts

Oracle In-Focus Series

by Robin Schumacher

published by Rampant TechPress 2003

Summary

Overall review: OK.

Target audience: Anyone who needs a quick overview of the causes of performance issues within the Oracle database.

Would you recommend to others? Yes, if you are new to performance tuning. Experienced DBAs will find it lacking.

Who will get the most from this book? Readers who have a general knowledge of the components of the Oracle database but haven't done performance tuning before.

Is this book platform specific? No, the book covers the specifics of the Oracle database and doesn't cover any operating systems issues. Note that this book does cover 9i features.

Why was this book read/chosen for review? I needed to understand the details of shared pool tuning, and I was looking for SQL scripts that would help me examine the relevant system tables.

Overall Review

The Technique You Will Learn

The back cover of the book promises to focus on how to quickly troubleshoot and correct Oracle performance problems. The discussion and scripts in the book do support this claim. However, the subject of Oracle tuning is much broader than what this book covers. Many readers may find that their database, and the environment in which real-world applications live, will have performance issues that aren't addressed by this book. The performance

troubleshooting approach described in this book relies on using SQL scripts exclusively. This will work for many—but not all—database performance issues. The more complex the environment, the less likely it is that this approach will provide a comprehensive solution. There are many performance issues that will affect overall application performance that aren't visible from within the database. Using only SQL scripts means your view of what's going on is going to be limited.

To make best use of the approach presented, you have to have already determined—from some other process—that the database performance is the primary performance issue for your application. For example, if the user's client machine is the primary reason for slow application response time, the database performance issues that can be detected and resolved by this approach may be irrelevant.

Chapter by Chapter

Chapter 1 – Accurately Measuring Performance – is a generic discussion of how to assess performance and what the various factors are that affect the end-user perception of overall application system performance.

Chapter 2 – Performance Methodologies – reviews the use of ratios, wait events, and looking for bottlenecks in the system. The comments on the limitations of each of these methodologies are valuable and should be given more weight in the mainstream DBA community. For example, it can be difficult to gauge the extent of sorts to disk from wait state analysis alone, while a ratio of in-memory sorts to disk sorts can be very useful. The author goes on to clearly explain the benefits and limitations of both ratio and wait based analysis.

Chapter 3 – Foundational Flaws – discusses the impact of physical database design errors. This is refreshing, as it doesn't get discussed much in the general DBA press. For those who haven't done much tuning work before, this foundation is needed. It is easy to spend lots of time tuning SQL when most of the problem is in the physical layout of the database. The author points out that application tuning can't be successful unless the physical layout of the database is done properly first.

Chapter 4 – Optimizing Storage – covers disk layout and fragmentation issues. The author expresses concern that bigger disks, which means fewer physical drives for a given database size, will cause performance problems. I haven't seen this in practice, so I think this would need to be supported with more evidence. The discussion of fragmentation is good, and provides a detailed example demonstrating that fragmentation does not cause performance problems.

Chapter 5 – Maximizing Memory – reviews the value (yes, the value) of buffer cache hit ratios, issues relating to the shared pool and sorting. A number of features new to Oracle 9i are covered. The author provides a good discussion of the use and misuse of the buffer cache hit ratio. This leads into a discussion of the new 9i feature `db_cache_advice` that provides some insight into how a larger buffer cache might help performance.

Chapter 6 – I/O Hotspots – covers how to find database objects, data files, and users that are causing the most I/O, chained rows and rollback issues. Scripts are presented to determine the I/O activity of the overall database, ratios for physical to logical I/O, wait events relating to I/O, how to see how objects are being accessed (table scan, row fetch), as well as the user sessions that are currently using the most I/O.

Chapter 7 – Workload Analysis Part 1, Problem Sessions – oddly, brings up security at the beginning of the chapter. The general issue raised is valid, but security should have been included in the high-level discussion of performance issues in Chapter 1. This section covers how to find users, processes, and individual SQL statements that are consuming the most resources. A personal favorite of mine is the script that will identify any SQL that is currently in the cache that is using a Cartesian join. I haven't seen this before, and it is a good way to find problematic SQL.

Chapter 8 – Workload Analysis Part 2, Problem SQL – a good high level review of how to find and fix problem SQL statements.

What Was Good

By using the scripts in the book, you can look into many Oracle database performance problems without installing any other tools. These scripts only require a SQL*Plus connection and the privileges needed to select from the system tables. For someone who can't get any more access, or someone who doesn't want to install any further tools such as STATSPACK or any of the many third party products available for Oracle tuning, this approach is appealing. It is the simplest, lowest cost and most universal way to approach Oracle database tuning.

The book was also inexpensive (relative to other Oracle books), and doesn't try to cover everything. As long as the reader understands what the book does and doesn't cover, the value is very good.

A fast read—some might say it's too short, but that can be viewed as a plus. This book doesn't waste your time with huge reprints of the Oracle manuals, like many tuning books do.

Small, targeted books like this one are to be encouraged, with the understanding that it may take more effort, and more editorial oversight, to make this work well. To be successful, a quick read needs to be very carefully targeted to deliver the best value to the reader.

What Could Be Better

The process described in the book is presented as Oracle performance troubleshooting, and covers perfor-

mance issues within the database. If you are sure the performance issue is in the Oracle database, that's fine. However, there are many performance issues that involve the operating system, network, client machine, etc. I think this should be pointed out more clearly in the early sections of the book.

The book doesn't mention the tools available to help find and resolve Oracle database performance issues. Specifically, I refer to STATSPACK and 10046 tracing. I have made clear my support of short, targeted books, and I don't want this book to expand to try to cover all possible database performance tools. However, short sections describing these tools would be good since these tools are included with the Oracle database software. Since I think this book is best for people who haven't done a lot of Oracle tuning, describing the existence and value of the tools that come with Oracle would be of great benefit.

A Review Relative to Tuning Books/Methods in General?

There are many sources for performance tuning methodologies. Because of the importance and popularity of this material, I like to look at any given approach with the following questions (along with my answers specific to this book):

How to approach tuning?

— Assumes the database is the performance issue, uses SQL scripts exclusively, doesn't use other tools available within Oracle (STATSPACK, 10046 trace)

What is the value of tuning?

— Not addressed

How to know when you are done?

— Not addressed

Can your users understand what you are doing?

— Yes, absolutely, process is clear

Is the process well documented?

— Yes, very much so

Is the process repeatable?

— Yes, definitely

Editorial Details

Normally, the physical layout and presentation of an Oracle book doesn't get my attention. Unfortunately, this one did. While the following comments don't detract from

**Don't miss our
Fall Conference!**
**November 4, 2004, at the
Computer History Museum
in Mountain View, CA!**

the value of the technical material in the book, the book itself appears to have been hastily prepared. It would be good if simple things had been done better. Each chapter should have a title, if only to show that the author and editor have a theme for the chapter and to show what the intended flow of the book was. (I know what it's like to have very good editors; they make you write better because you know someone, a very demanding someone, is always watching!). Further, it would be good if each page of the book showed which chapter it was part of. This book simply repeats the book title at the bottom of each page.

On page 13, halfway down the page, there is a mysterious box that says, "The code depot key is rudy". What is this about? I have no clue what this is telling me. Tight editing could have prevented this.

As mentioned previously, one of the specific reasons I spent my own money on this book was for the scripts, but I don't want to even try to manually type them into my system from the book. The book comes with a CD, but the CD doesn't have the scripts. On page 1 of the book we are told about the Online Code Depot, the URL is given and we are told that this is where to go for the sample tests and answers. I don't see any sample tests in this book, and there is no mention of whether the scripts are available from any source. I assumed that the reference to "sample tests" really meant the scripts in the book, so I tried the URL. But, you have to have a userid and password for the URL. I had to email the publisher asking for the userid and password (and the password isn't "rudy").

They replied quickly, but this process could have been much simpler.

Conclusion

Overall, this book could be ideal for developers that are supporting their own development databases. In this environment, the developers will have full access to the database and the database machine and may not have a full time DBA to monitor database performance. With the efficient presentation of the performance issues, and the scripts, it would be easy to monitor and resolve performance issues as the developers modify their code. The book is quick read which makes it a good value for anyone that wants a quick, high-level view of the primary performance issues within the Oracle database. ▲

About the Book Author

Robin Schumacher is vice-president of Product Management for Embarcadero Technologies, Inc., a leading supplier of database software tools. Robin has over fourteen years experience in database administration, development, monitoring, and tuning with Oracle, DB2, Teradata, Sybase, and Microsoft SQL Server. He has authored countless performance-related articles for many database-centric magazines as well as serving as a database software reviewer and feature writer for the likes of Intelligent Enterprise, eWeek, DM Review, and others.

(See page 27 for information about the book's reviewer, Brian Hitchcock.)

A Little NoCOUG History

In honor of our upcoming visit to the Computer History Museum for our Fall Conference on November 4, we thought we'd publish a few facts about NoCOUG's history. It's interesting to see how we've grown.

- NoCOUG started in the late 1980s.
- Our Board member with the most knowledge of NoCOUG history is Joel Rosingana. He joined NoCOUG in about 1991 and has been on the Board of Directors since about 1993. He has held many positions, including president (more than once). He has a little history in his office, too. As many of his friends know, he has Oracle v6 on DOS. But don't worry, that's for nostalgic purposes. He's also running Oracle 9i.
- Back in the early days, 60 attendees at a conference was a good turnout. At the Summer NoCOUG Conference in August 2004, we had over 250 attendees.
- In 1991, individual membership dues were \$50. For 2005, they are \$70.
- At the Spring Conference in 1991, there was one track and four technical sessions, and the meeting ran from 9:00 a.m. to 2:30 p.m. At the Spring Conference in 2004, there were three tracks and 12 technical sessions, and the meeting ran from 8:00 a.m. to 4:45 p.m.
- NoCOUG's bank balance as of September 30, 1991, was \$6,592.80. On September 30, 2004, NoCOUG's balance was \$60,908.56. ▲