

Official Publication of the Northern California Oracle Users Group

# NoCOUG

J O U R N A L

Vol. 23, No. 2 · MAY 2009

\$15

## Let Knowledge Spring Forth at NoCOUG

### Fresh Perspectives

*An interview with Karen Morton.*

*See page 4.*

### ASM Test Environment

*An excerpt from a new book.*

*See page 9.*

### SQL Corner

*Graphical query execution plans.*

*See page 16.*

*Much more inside . . .*

# Troubleshooting Oracle Performance

A Book Review by Raghav Vinjamuri

## Details

**Author:** Christian Antognini

**ISBN:** 978-1-59059-917-4


**Pages:** 586

**Year of Publication:** 2008

**Edition:** 1

**Price:** \$59.95

**Publisher:** Apress



## Summary

**Overall review:** This book explains what you should not forget, from a performance point of view, the next time you are developing an application. This is not a book for beginners.

**Target audience:** Oracle application developers and designers.

**Would you recommend to others:** Yes.

**Who will get the most from this book:** Intermediate and experienced developers and designers with a working knowledge of the Oracle database engine and proficiency with SQL.

**Is this book platform specific:** No.

**Why did I obtain this book:** It was recommended to me.

## Overall Review

In the first page of the first chapter, the author summarizes the contents of the book in the following words: “*Real development teams quite often forget about performance, until at least performance problems arise. At this point, it may be too late. Therefore, in the following sections, I’ll cover what you should not forget, from a performance point of view, the next time you are developing an application.*” A round of applause for the author for recognizing that performance must be kept in mind at every step of the Software Development Life Cycle (SDLC).

The book is divided into four parts. Part I covers the basics. It introduces the subject and covers key concepts. Part II introduces a methodology for identifying performance problems and the tools that can be used to do so. Part II focuses on the query optimizer, and Part IV focuses on the performance

features of the database engine.

This is not a book for beginners. For example, in Chapter 2, the life cycle of a cursor is discussed in great depth without first explaining what a cursor is. Other terms such as implied cursors, parent cursors, and child cursors are also used without explanation. Nevertheless, it was a joy to read this book.

## Chapter 1—Performance Problems

The book opens with a discussion of response time and its components: service time and wait time. I could not have asked for a better start. Perhaps the most important clause in this chapter is “set a quantifiable target for the optimization”—it cannot be emphasized enough.

## Chapter 2—Key Concepts

This chapter starts with simple topics such as selectivity and cardinality and then progresses into more advanced topics such as cursors, bind-variable peeking, and extended cursor sharing. A short introduction to the concept of a cursor would have been in order. The discussion on the life cycle of a cursor is an enthralling read.

## Chapter 3—Identifying Performance Problems

Chapter 3 is the sole chapter in Part II of the book and coincidentally is the longest chapter in the book. Here the focus is on describing a methodology that can be used to find out where and how time is spent. It is an excellent compendium of the various tools, tips, and suggestions available to a developer to identify problem SQL statements. There is a decent discussion of instrumentation and profiling. About ten pages are devoted to commercial products such as PerformaSure and Jprobe. Finally, there is a 40-page discussion on `sql_trace`, `tkprof`, and the author’s own tool, TVD\$XTAT (Trivadis Extended Tracefile Analysis Tool). Personally, I would have liked to see discussion of non-SQL problems; for example, a server with an inadequately sized shared pool.

## Chapter 4—System and Object Statistics

This chapter kicks off Part III and is the second-longest chap-

*“Real development teams quite often forget about performance, until at least performance problems arise. At this point, it may be too late. Therefore, in the following sections, I’ll cover what you should not forget, from a performance point of view, the next time you are developing an application.”*

ter in the book. The author states that there are two major causes of performance problems in an Oracle database: not designing applications for performance and poor configuration of the query optimizer. The chapter starts with an overview of `dbms_stats` and explains the differences between system statistics and object statistics. The discussion of “pending statistics”—an Oracle Database 11g feature—caught my attention.

---

### Chapter 5—Configuring the Query Optimizer

---

The chapter identifies how to configure initialization parameters that affect the query optimizer, including `optimizer_mode`, `optimizer_features_enable`, `db_file_multiblock_read_count`, `optimizer_dynamic_sampling`, `optimizer_index_cost_adj`, `optimizer_index_caching`, `optimizer_secure_view_merging`, `workarea_size_policy`, `pga_aggregate_target`, `sort_area_size`, `sort_area_retained_size`, `hash_area_size`, and `bitmap_merge_area_size`.

---

### Chapter 6—Execution Plans

---

This chapter discusses how to obtain and interpret execution plans and how to recognize inefficient execution plans. Utilities and interfaces such as EXPLAIN PLAN, dynamic performance views, various tracing tools, and the `dbms_xplan` package are discussed. I liked the treatment of the `dbms_xplan.display`; the `dbms_xplan.display_cursor`; and the `dbms_xplan.display_awr` functions for returning execution plans stored in the `plan_table`, the `library_cache`, and the AWR.

---

### Chapter 7—SQL Tuning Techniques

---

This chapter surveys SQL tuning techniques such as hints, SQL profiles, and stored outlines. I liked the discussion of the pitfalls and fallacies of each technique.

---

### Chapter 8—Parsing

---

This chapter kicks off Part IV with a detailed discussion of query parsing. It explains quick parses and long parses and what to do if you encounter excessive parsing.

---

### Chapter 9—Optimizing Data Access

---

This chapter explains how indexing and partitioning can be used to improve access to the data.

---

### Chapter 10—Optimizing Joins

---

This chapter discusses the different join algorithms—Nested Loop Joins, Merge Joins, and Hash Joins—and their applicability. This is a mostly theoretical chapter.

---

### Chapter 11—Beyond Data Access and Join Optimization

---

This chapter discusses an assortment of topics such as materialized views, result caching, parallel processing, direct-path insert, and pre-fetching. As in Chapter 7, the author discusses the pitfalls and fallacies of each technique.

---

### Chapter 12—Optimizing the Physical Design

---

This chapter discusses optimal column order, best practices in datatype selection, row migration and row chaining, block contention, and data compression, among other physical design issues. ▲

---

*Raghav Vinjamuri has a strong background in Oracle SQL and PL/SQL development and tuning, performance tuning and scalability issues, Linux device drivers, Unix kernels, and shell programming. He is a certified trainer and an Oracle OCP-certified DBA.*