

# Foreword

**A**rthur C. Clarke once wrote that “any sufficiently advanced technology is indistinguishable from magic.” I believe that observation to be entirely accurate. Someone else later observed that “any technologist with sufficient knowledge is indistinguishable from a magician.” With that in mind, what you have in your hands right now is a book on magic.

However, for as long as I’ve known the author of this book, Jonathan Lewis (some 11 years according to my research on Google going back over the newsgroup archives), he has never been content to accept “magic.” He wants to know *why* something happens the way it does. So, fundamentally, his book is all about *understanding*:

*understand v. understood, (-std) understanding, understands*

1. *To perceive and comprehend the nature and significance of*
2. *To know thoroughly by close contact or long experience with*

More precisely, it is all about understanding the Oracle cost based optimizer (CBO), how it works, and why it does what it does. Jonathan conveys to us his understanding of the Oracle CBO through practice and example, and with this understanding, with this knowledge, new options and solutions become available.

Put simply, the Oracle CBO is a mathematical model; you feed it inputs (queries, statistics), and it produces outputs (query plans). In order to use the CBO successfully, it is critical that you understand what these inputs are and how the CBO uses them. Consider the following question, one that I’ve tried to answer many times: What is the best way to collect statistics, what statistics should I gather? Seems simple enough, very straightforward—there should be an answer and there is, but it is no “one size fits all” answer. It depends on your circumstances, your data distribution, your queries, the type of system you have (transactional, data warehouse)—a whole host of factors—and it is only through an understanding of how the CBO works and of how these factors affect the CBO that you’ll be able to answer the question for yourself.

My favorite chapter in this book is Chapter 7, which provides an excellent discussion on what histograms do, how the CBO uses them, and some of the myths that surround them (they are a frequently misunderstood input to the CBO). One of the reasons it is my favorite chapter is because I still remember *hearing* this chapter for the first time (not reading it, hearing it). It was about three years ago, at the NoCOUG (Northern California Oracle User Group) meeting. I attended Jonathan’s *Histograms* session and, for the first time, felt that I truly *understood* how histograms worked in Oracle. Jonathan provided practical information, of immediate day-to-day use, and as a result I was in a much better position to answer the question just posed. In this book, you will find many such topics described in that same practical, immediately usable way.

The insights that Jonathan provides into the workings of the cost based optimizer will make a DBA a better designer and a developer a better SQL coder. Both groups will become better troubleshooters as a result. Jonathan’s meticulous and abundant examples make the complex workings of the CBO understandable for the rest of us.