

Book Review: Rob Verschoor's New Cookbook

Published by Sypron Publications, 2003, ISBN 90-806117-2-7

Reviewed by Anthony Mandic

Anyone familiar with Rob Verschoor from the Sybase newsgroups and mailing lists knows that he has a talent for pulling rabbits out of hats. Ever now and again he will present a solution to a curly problem that astounds by its simplicity and elegance. He has frequently presented such tips and tricks at TechWave conferences and has now followed up with a book in a similar vein.

Rob has a new book out called *Tips, Tricks, and Recipes for Sybase ASE* (see page 13). At 416 pages with 18 chapters and appendices, Rob once again proves his mettle as an author, with diverse solutions to many of the problems facing any developer of Transact-SQL and related scripting languages. The first chapter explores some of the history of Sybase SQL Server and ASE before examining the various versions and their associated nomenclature, and the many ways of determining which version you are working with. Some ASE internals are covered throughout the book as well. This is useful in understanding why ASE functions the way it does.

Chapter 2 examines ASE's associated files—in particular, the error log and some of the things you can do with it using T-SQL. Chapter 3 takes an in-depth look at killing spids and terminating processes. Don't be fooled into thinking this is a short chapter. Far from it. Its size shows the lengths Rob goes to in providing detailed information.

Chapter 4 examines methods for dealing with duplicate data in a whole raft of situations. Readers should find something here for any situation they are likely to encounter. Chapter 5 is somewhat related, in that it looks at what you can do with the *ignore_dup_key* index option. Once again, the amount of detail doesn't make for a lightweight chapter. Chapter 6 examines the complex issue of missing data and how to go about finding it (if this sounds absurd, reading the chapter will help clarify it).

Chapters 7 and 8 provide quite a few recipes for generating primary keys. Although this is another understatement, the chapters are quite exhaustive. The former chapter looks at sequential primary keys and the latter non-sequential.

Chapter 9 examines all the things you can do with the **update** statement.

I should point out that Rob also examines the performance aspects of the various examples he provides, which invariably is useful for developers to know. Chapter 10 deals with the **execute immediate** statement and dynamic T-SQL.

Chapter 11, meanwhile, deals with identity columns in some detail. This includes how they work, the issues with identity gaps and how to deal with them or avoid them, identity burning set factor, and dump and load problems related to identity gaps.

In Chapter 12, Rob looks at transactions and transaction log issues, spids and some related internals. Chapter 13 presents various ASE internals in more detail and Chapter 14 covers some tools that could prove useful. The appendices examine recent changes to ASE, getting diagnostic information out, identity columns prior to ASE 12, and the electronic supplement to the book.

The book also contains a lot of T-SQL code which form the recipes. Fortunately, the code is obtainable from Rob's website. The book is available at www.sypron.nl/ttr as of July 2003. ■