

Database Benchmarking and Stress Testing

**An Evidence-Based Approach
to Decisions on Architecture
and Technology**

Bert Scalzo

Apress®

Database Benchmarking and Stress Testing

Bert Scalzo
Flower Mound, TX, USA

ISBN-13 (pbk): 978-1-4842-4007-6
<https://doi.org/10.1007/978-1-4842-4008-3>

ISBN-13 (electronic): 978-1-4842-4008-3

Library of Congress Control Number: 2018960192

Copyright © 2018 by Bert Scalzo

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr
Acquisitions Editor: Jonathan Gennick
Development Editor: Laura Berendson
Coordinating Editor: Jill Balzano

Cover designed by eStudioCalamar

Cover image designed by Freepik (www.freepik.com)

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail rights@apress.com, or visit <http://www.apress.com/rights-permissions>.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub via the book's product page, located at www.apress.com/9781484240076. For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

*To my wife, Susan, who has put up with me for 30+ years
and our four legged children over the years:
Ziggy, Max, and Dexter.*



Table of Contents

About the Author	xi
About the Technical Reviewer	xiii
Acknowledgments	xv
Introduction	xvii
Chapter 1: Benchmarking Basics	1
What Is a Database Stress Test?	2
What Is a Database Benchmark?	3
What Is Workload Capture/Replay?	8
Why These Techniques Are Useful.....	9
When These Techniques Are Valuable.....	10
How to Choose Between These Options	11
Where We Go from Here.....	12
Chapter 2: Industry Standard Benchmarks	15
Most Popular Benchmarks.....	16
TPC-C.....	16
TPC-E.....	21
TPC-H	25
TPC-DS	30
Less Popular Benchmarks	38
TPC-DI.....	38
TPC-VMS.....	40

TABLE OF CONTENTS

Wisconsin	41
AS ³ AP	41
Obsolete Benchmarks	42
TPC-A.....	42
TPC-B.....	43
TPC-D	43
TPC-R.....	44
TPC-W.....	44
Summary	45
Chapter 3: Benchmarking Tools.....	47
Storage Benchmarking	48
WINSAT	49
HDPARM	50
SQLIO	50
DISKSPD	51
Database Simulation	52
SQLIOSTRESS	52
SQLIOSIM.....	53
Orion.....	54
Calibrate IO	56
Database Benchmarking.....	59
HammerDB	59
Benchmark Factory	84
Swingbench.....	91
SLOB.....	96
Summary.....	100

Chapter 4: Benchmarking Preparation	101
Benchmark Preparation	102
Database Preparation.....	106
Benchmarking Time Line	108
Benchmark Sizing	110
Workload Capture	111
Distorting Factors.....	112
Summary.....	114
Chapter 5: Benchmarking Mistakes	115
Keep References Handy	116
All the Required Tools	117
No Big Red Easy Button	119
Proper Tool Deployment.....	123
Unexpected Human Factors	125
Validate Our Direction.....	125
Self-Fulfilling Prophecy	126
Top Ten Benchmarking Misconceptions.....	126
#1: I'm using a tool like Quest's Benchmark Factory, so that's all I need....	127
#2: I have an expensive SAN, so I don't need to configure anything special for IO.....	128
#3: I can just use the default operating system configuration right out of the box	129
#4: I can just use the default database setup/configuration right out of the box	129
#5: Tools like Benchmark Factory will create optimally designed database objects for my hardware.....	130

TABLE OF CONTENTS

#6: Tools like Benchmark Factory will automatically monitor, tune, and optimize all my hardware, operating system, and database configuration parameters	130
#7: I don't need a DBA to perform benchmark tests – anyone technical can do it	131
#8: I can very easily compare database vendors on the same hardware platform.....	131
#9: Transactions per Second (i.e., TPS) are what matter most in benchmarking.....	132
#10: We can spend just a few days and benchmark everything we're interested in	133
Summary.....	134
Chapter 6: Benchmarking Hardware Options	135
Draw a Picture	136
CPU and Memory	138
Spinning Disks	140
New Storage Technologies.....	142
Specialized Appliances	144
Summary.....	145
Chapter 7: Benchmarking Software Options	147
Draw a Picture	148
Partitioning Larger Objects	152
Advanced Data Compression	155
SSD to Extend Memory	157
Pinning Tables in Memory.....	159
Columnstore vs. Rowstore	161
New “In-Memory” Tables	164
Summary.....	167

Chapter 8: Benchmarking for Consolidation.....	169
Consolidation Approaches.....	170
Database Coexistence.....	174
Unforeseen Issues.....	176
Server Distribution.....	177
Mixing Benchmarks.....	178
Summary.....	183
Chapter 9: Benchmarking for Virtualization.....	185
Server BIOS.....	186
VM Configuration.....	187
CPU.....	187
Memory.....	190
Network.....	192
Disk IO.....	194
Monitoring Issues.....	196
Summary.....	197
Chapter 10: Benchmarking for Public Cloud.....	199
Image Right Sizing.....	200
Deployment Architecture.....	202
Summary.....	206
Chapter 11: Workload Capture and Replay.....	207
Double Everything.....	208
Capture/Replay Tools.....	210
Capture Duration.....	211
Replay Architecture.....	213

TABLE OF CONTENTS

Interpreting Results 215

Capture/Replay Example..... 216

Summary..... 220

Chapter 12: Database Stress Testing.....221

Raw Stress Testing..... 221

DBBENCHMARK Tool 223

Summary..... 234

Index.....235

About the Author

Bert Scalzo is an Oracle ACE, blogger, author, speaker, and database technology consultant. He has a BS, MS, and PhD in computer science; an MBA; and has worked for over 30 years with all major relational databases, including Oracle, SQL Server, DB2 LUW, Sybase, MySQL, and PostgreSQL. Moreover Bert has also worked for several of those database vendors. He has been a key contributor for many popular database tools used by millions of people worldwide, including TOAD, Toad Data Modeler, ERwin, ER/Studio, DBArtisan, Aqua Data Studio, and Benchmark Factory. In addition, Bert has presented at numerous database conferences and user groups, including SQL Saturday, SQL PAAS, Oracle Open World, DOUG, ODTUG, IOUG, OAUG, RMOUG, and many others. His areas of interest include data modeling, database benchmarking, database tuning and optimization, “star schema” data warehouses, Linux®, and VMware®. Bert has written for Oracle Technology Network (OTN), *Oracle Magazine*, *Oracle Informant*, *PC Week* (eWeek), *Dell Power Solutions Magazine*, *The LINUX Journal*, LINUX.com, Oracle FAQ, and Toad World. Moreover Bert has written an extensive collection of books on database topics, focusing mainly around TOAD, data warehousing, database benchmarking, and basic introductions to mainstream databases.

About the Technical Reviewer



Arup Nanda has been an Oracle database administrator (DBA) since 1993, dealing with everything from modeling to security, and he has a lot of gray hairs to prove it. He has coauthored 5 books, written 500+ published articles, presented 300+ sessions, delivered training sessions in 22 countries, and actively blogs at arup.blogspot.com. He is an Oracle ACE Director, a member of Oak Table Network, an editor for *SELECT Journal* (the IOUG publication), and a member of the Board

for Exadata SIG. Oracle awarded him the DBA of the Year in 2003 and Architect of the Year in 2012. He lives in Danbury, CT, with his wife Anu and son Anish.

Acknowledgments

I have been doing database work, including benchmarking, for over three decades. Much of what I've learned has come from working alongside other great people who were very willing to share their expertise. So rather than list a few people and risk missing anyone, let me just say thanks to all those people “along the way” who’ve helped me to learn so much.

I also owe a debt of gratitude to the following companies for excellent opportunities and experience:

- IDERA
- Quest Software
- Embarcadero Technologies
- Logic Works
- Dell
- IBM
- EDS
- Citicorp
- Oracle Corporation
- Nationwide Insurance
- Battelle Memorial Institute

Introduction

Database Benchmarking and Stress Testing is for database administrators and professionals responsible for advising on any and all database architectural decisions ultimately affecting overall database performance. This book provides an empirical method for answering “what if” questions about database performance, helping database administrators make sound architectural decisions in a fast-changing landscape of virtualized servers and container-based solutions. Today’s database administrators face numerous such questions:

- What if we consolidate databases using multitenant features?
- What if we virtualize database servers as Docker containers?
- What if we deploy the latest in NVMe flash disks to speed up IO access?
- Do features such as compression, partitioning, and in-memory OLTP earn back their price?
- What if we move our databases to the cloud?

As an administrator, do you know the answers or even how to test the assumptions? You should know how to test and find answers to these questions, and doing so is what this book is about. *Database Benchmarking and Stress Testing* introduces you to database benchmarking using industry-standard test suites such as the TCP series of benchmarks, which are the same benchmarks that vendors rely upon. You’ll learn to run these industry-standard benchmarks and collect results to use in answering

INTRODUCTION

questions about the performance impact of architectural changes, technology changes, even down to the brand of database software. You'll learn to measure performance and predict the specific impact of changes to your environment. You'll learn the limitations of the benchmarks and the crucial difference between benchmarking and workload capture/reply.

Database Benchmarking and Stress Testing is about creating empirical evidence in support of business and technology decisions. It's about not guessing when you should be measuring. Empirical testing is scientific testing that delivers measurable results. Begin with a hypothesis about the impact of a possible architecture or technology change. Then run the appropriate benchmarks to gather data and predict whether the change you're exploring will be beneficial, and by what order of magnitude. Stop guessing. Start measuring. Let *Database Benchmarking and Stress Testing* show the way.