



ORACLE®

ORACLE PRESS™ — EXCLUSIVELY FROM McGRAW-HILL/OSBORNE

# ORACLE DATABASE 10g

# Linux Administration

Deploy and Optimize Oracle Database 10g in a Linux Environment



ORIGINAL • AUTHENTIC

***Oracle Press™***

ONLY FROM OSBORNE

**EDWARD WHALEN**

Founder and Chief Technologist of  
Performance Tuning Corporation

MCGRAW HILL COMPANIES  
FOR DONATION ONLY  
NOT FOR RESALE  
12ACA07

FOR DONATION ONLY  
NOT FOR RESALE

# JOIN THE **ORACLE PRESS** **COMMUNITY TODAY!**

## Get these member benefits:

- Get "sneak peeks" at Oracle Press titles - not yet published
- Receive special promotional offers
- Read "New Book Alerts" to learn about the latest Oracle Press releases
- Hear from the editor as to what's coming and why

GIFT

Q1

A FOUNDATION

AND MUCH MORE!

Y CHÂU Á  
BÁN LẠI

JOIN TODAY AT [WWW.ORACLEPRESSBOOKS.COM](http://WWW.ORACLEPRESSBOOKS.COM)



ORIGINAL • AUTHENTIC

**Oracle Press**

ONLY FROM OSBORNE

**ORACLE®**

*Oracle Press™*

# Oracle Database 10g Linux Administration

Edward Whalen



**GIFT OF THE ASIA FOUNDATION  
NOT FOR RE-SALE**

**QUÀ TẶNG CỦA QUỸ CHÂU Á  
KHÔNG ĐƯỢC BÁN LẠI**

**McGraw-Hill/Osborne**

New York Chicago San Francisco  
Lisbon London Madrid Mexico City Milan  
New Delhi San Juan Seoul Singapore Sydney Toronto

**McGraw-Hill/Osborne**  
2100 Powell Street, 10th Floor  
Emeryville, California 94608  
U.S.A.

To arrange bulk purchase discounts for sales promotions, premiums, or fund-raisers, please contact **McGraw-Hill/Osborne** at the above address.

## **Oracle Database 10g Linux Administration**

Copyright © 2005 by The McGraw-Hill Companies, Inc. (Publisher). All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of Publisher.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.  
Screen displays of copyrighted Oracle software programs have been reproduced herein with the permission of Oracle Corporation and/or its affiliates.

1234567890 FGR FGR 0198765

ISBN 0-07-223053-3

**Acquisitions Editor**

Lisa McClain

**Project Editor**

Carolyn Welch

**Contributing Authors**

Robert Bryan Thomas

Fraser Talbot

Dan Hotka

**Technical Editors**

Michael New

James Petts

Greg Marsden

Joel Becker

**Copy Editor**

Bob Campbell

**Proofreader**

Susie Elkind

**Indexer**

Claire Splan

**Composition**

Apollo Publishing Services

**Cover Series Design**

Damore Johann Design, Inc.

**Series Design**

Jani Beckwith, Peter F. Hancik

This book was composed with Corel VENTURA™ Publisher.

Information has been obtained by Publisher from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Publisher, or others, Publisher does not guarantee to the accuracy, adequacy, or completeness of any information included in this work and is not responsible for any errors or omissions or the results obtained from the use of such information.

Oracle Corporation does not make any representations or warranties as to the accuracy, adequacy, or completeness of any information contained in this Work, and is not responsible for any errors or omissions.

## About the Author

**Edward Whalen** is the founder of Performance Tuning Corporation ([www.perftuning.com](http://www.perftuning.com)), a consulting company specializing in database performance, administration, and backup/recovery solutions. Prior to starting Performance Tuning Corp., Mr. Whalen worked at Compaq Computer Corporation as an OS developer, and then as a database performance engineer. He has extensive experience in database system design and tuning for optimal performance. His career has consisted of hardware, OS, and database development projects for many different companies. Mr. Whalen has written three other books on the Oracle RDBMS and four books on MS SQL Server. He has also worked on numerous benchmarks and performance tuning projects with both Oracle and MS SQL Server. Mr. Whalen is recognized as a leader in database performance tuning and optimization. In addition to his experience in Oracle on many different OS platforms, such as Solaris, AIX, HP-UX, and Windows, he has recently spent significant time working with Oracle on Linux, including the deployment of some of the first RAC clusters on Linux. He works closely with both Red Hat and Novell/SUSE with many different flavors of Linux.

## About the Contributing Authors

**Robert Bryan Thomas** is a Senior Performance Consultant at Performance Tuning Corporation. He has a diverse background in database performance, Oracle RAC, data warehousing, and SAN storage. He currently specializes in Oracle10g RAC on multiple OS platforms. Robert Thomas graduated from Texas A&M University with a degree in Computer Science.

**Fraser Talbot** is a Senior Performance Consultant at Performance Tuning Corporation. His background includes several years of database administration, exclusively Oracle and Oracle E-Business Suite. He and his wife, Virginia, and three daughters, Margaret, Susan, and Elizabeth, reside in Houston, Texas.

**Dan Hotka** is a Training Specialist, who has over 27 years of experience in the computer industry and over 22 years of experience with Oracle products. He is an internationally recognized Oracle expert with Oracle experience dating back to the Oracle V4.0 days. He is also the author of *Oracle9i Development by Example* and *Oracle8i from Scratch* by Que and has co-authored six other popular books. He is frequently published in Oracle trade journals, and speaks regularly at Oracle conferences and user groups around the world. Visit his website at [www.DanHotka.com](http://www.DanHotka.com). Dan can also be reached at [dhotka@earthlink.net](mailto:dhotka@earthlink.net).



# Contents

ACKNOWLEDGMENTS .....	xv
INTRODUCTION .....	xvii

## PART I Oracle on Linux Overview

<b>1 Linux Architecture .....</b>	<b>3</b>
Operating System Overview .....	5
The Monolithic Kernel .....	5
The Microkernel .....	6
The Hybrid Kernel .....	7
Linux Overview .....	8
The Linux Directory Structure .....	8
Linux Utilities and Directories .....	9
The Development Environment .....	13
The Linux User Interface .....	17
X Windows .....	17
GUI Environments .....	18
Linux Web Server Utilities .....	19
The Apache Web Server .....	19
CGI Programming in Linux .....	19
The Linux Boot Process .....	20
Linux Source Code .....	20
Rebuilding the Linux Kernel .....	21
Processes and Threads .....	23
What Is a Process? .....	23
Oracle's Use of Processes .....	23
What Is a Thread? .....	23
Devices .....	24
Character Devices .....	24

Block Devices .....	25
Logical Volume Managers .....	25
Automatic Storage Management .....	25
Filesystems .....	26
ext3 .....	26
OCFS .....	26
Raw Devices .....	27
Summary .....	<b>29</b>
<b>2 Oracle Architecture .....</b>	<b>30</b>
The Oracle Database Server .....	31
The Oracle Database .....	31
Physical Data Structures .....	35
Logical Data Structures .....	54
The Oracle Instance .....	54
The Instance Defined .....	54
Components of the Instance .....	58
User Processes .....	59
Server Processes .....	59
PGA Memory .....	59
MTS vs. Dedicated Server Processes .....	59
Oracle 10g New Features .....	60
Performance and Scalability Features .....	60
Clustering .....	61
Availability .....	61
Manageability .....	61
Business Intelligence .....	62
Application Development .....	62
Linux 64-Bit Support .....	63
Grid Computing .....	63
Summary .....	63
<b>3 Transaction Management and the Oracle Schema Objects .....</b>	<b>65</b>
Oracle Functionality .....	66
Reading from the Database .....	66
Concurrency .....	67
Writing to the Database .....	67
Locking .....	67
Read Consistency .....	68
Data Integrity .....	68
The Redo Log .....	69
Undo .....	76
Introduction to Transaction Management .....	79
Transactions .....	79

Introduction to the Oracle Schema Objects .....	83
Tables .....	84
Indexes .....	85
Views .....	86
Sequences .....	87
Clusters .....	87
Synonyms .....	87
Summary .....	87
 PART II	
<b>Deploying Oracle 10g on Linux</b>	
<b>4 Sizing Oracle 10g on Linux Systems .....</b>	<b>91</b>
An Introduction to Sizing .....	92
Sizing for Peak Utilization and Steady-State Utilization .....	92
Service-Level Agreements .....	93
System Resources, Capacity, and Queuing .....	95
CPU Performance and Queuing .....	96
I/O Performance and Queuing .....	97
Network Performance .....	98
Memory .....	98
Metrics .....	99
Oracle Sizing .....	101
Limitation Parameters .....	101
Resource Parameters .....	102
Hardware Sizing .....	104
CPU Sizing .....	104
Memory Sizing .....	104
I/O Subsystem Sizing .....	105
RAID Disk Subsystems .....	114
Hardware vs. Software RAID .....	114
Striping .....	115
Parity Overview .....	120
Capacity Planning .....	123
Summary .....	124
<b>5 Configuring the Linux System for Oracle .....</b>	<b>125</b>
Preinstallation Procedures for Linux .....	126
Minimum Requirements .....	126
Configuring the Filesystems .....	127
Choosing Linux Components .....	128
Post-Installation Procedures for Linux .....	131
Configuring Kernel Parameters .....	131

Adding Users and Groups .....	133
Configuring Shell Limits for the Oracle User .....	134
Creating Directory Structures .....	134
Setting Up the Oracle Environment .....	135
Installation Checklists .....	136
Installation Checklist for Oracle 10g R1 .....	137
Installation Checklist for Oracle 10g R2 .....	138
Summary .....	140
<b>6 Installing Oracle 10g on Linux .....</b>	<b>141</b>
Preinstallation Steps .....	142
Kernel Configuration Parameters .....	142
Storage Requirements .....	146
Device Names Might Change .....	149
Configuring the ASM Library Driver .....	149
Setting Up ASM Drives .....	150
The Installation Process .....	154
Staging the Install .....	154
Running the Installer .....	155
Postinstallation Steps .....	158
Downloading and Installing Patches .....	158
Configuration Steps .....	159
Testing and Validating .....	162
Additional Miscellaneous Configuration Steps .....	162
Summary .....	162
<b>PART III</b>	
<b>Oracle RAC on Linux</b>	
<b>7 Oracle Real Application Cluster Concepts .....</b>	<b>165</b>
What Is Clustering? .....	167
Cluster Terminology .....	167
Cluster Definition .....	169
RAC (Real Application Clusters) .....	171
Overview of RAC .....	171
Oracle RAC 10g Components .....	174
RAC Database Components .....	178
Oracle 10g RAC's Use of Shared Storage .....	181
Using NAS as Shared Storage .....	183
Summary .....	184
<b>8 Deploying Oracle Real Application Clusters on Linux .....</b>	<b>185</b>
Planning the RAC System .....	186
The RAC Components .....	186
Sizing the RAC System .....	189

Configuring Linux .....	190
Verifying Hardware .....	191
Verifying Network Requirements .....	192
Setting Up the Linux Kernel .....	193
Setting Up the Linux System .....	196
Configuring Linux Modules/Programs .....	198
Oracle Directories .....	198
Setting Environment Variables .....	200
Networking .....	200
Configuring RAC .....	203
Configuring the Cluster Ready Services (CRS) .....	204
Configuring the Interconnect .....	205
Configuring the Shared Disk .....	206
Testing and Managing the RAC .....	208
Oracle 10g RAC Help .....	213
Summary .....	213
<b>9 Administering Oracle and RAC on Linux .....</b>	<b>215</b>
Configuring the Linux Kernel .....	216
Configuring Shared Memory .....	216
Configuring Semaphores and Other Kernel Parameters .....	216
Monitoring and Changing Linux Configuration Parameters .....	217
Configuring the Network .....	218
Monitoring Linux .....	218
sar .....	219
vmstat .....	220
top .....	220
iostat .....	221
Upgrading the Linux Operating System .....	222
Starting Up the Instance .....	222
Database Startup .....	223
Database Mounted .....	224
Database Opened .....	224
Starting Up from the Enterprise Manager Database Control .....	225
Oracle Database 10g and CSSD .....	227
User Authentication .....	227
Database Authentication .....	227
Operating System Authentication .....	227
Network Authentication .....	227
Remote Authentication .....	228
Mid-Tier Authentication .....	228
Authentication by SSL .....	228
User Authorization .....	229
User Profiles .....	229
Privileges .....	231

Adding a User Account .....	233
Summary .....	234
<b>10 Tuning Oracle and RAC on Linux</b> .....	<b>235</b>
Performance Tuning Methodology .....	236
Step 1: Initial Assessment .....	237
Step 2: Monitor the System .....	238
Step 3: Analyze Results .....	239
Step 4: Create a Hypothesis .....	239
Step 5: Propose a Solution .....	239
Step 6: Implementation of Solution .....	240
Step 7: Test and Monitor .....	241
Further Analysis and Testing .....	241
Tuning Oracle on Linux .....	242
Application Tuning .....	243
Instance Tuning Basics .....	243
Monitoring the Oracle Instance .....	246
Tuning for Wait Events .....	253
Hardware Tuning .....	253
Monitoring the Linux System .....	254
Monitoring Linux with ps .....	255
Monitoring Linux Using vmstat .....	257
Monitoring Linux Using iostat .....	257
Monitoring Linux Using sar .....	258
Monitoring Linux Using Top .....	259
Using Large Memory Models .....	260
Configuring Linux for > 1.7GB (and <= 2.7GB) of SGA .....	260
Configuring Oracle for > 2.7GB SGA .....	261
Using Load Testing for Performance Validation .....	262
Why Is the System Being Tested? .....	263
Load-Testing Batch Jobs .....	264
Load-Testing Loading .....	265
Load-Testing OLTP Applications .....	265
Performance Metrics .....	267
Load Testing Tips .....	268
Summary .....	269

PART IV

**Administering and Tuning Oracle Database 10g on Linux**

<b>11 Backup and Recovery</b> .....	<b>273</b>
Backup Overview .....	274
Backup and Recovery Concepts .....	275
Instance Recovery .....	275
Data Recovery or Media Recovery .....	276

Online and Archived Redo Log File Recovery .....	276
How Transactions Work .....	278
RDBMS Functionality .....	279
Introduction to Backup .....	283
Archivelog Backup .....	283
Controlfile Backup .....	285
Full Backup .....	285
Tablespace and Datafile Backups .....	286
Incremental Backups .....	287
How to Back Up the Database .....	288
Recovery Manager (RMAN) .....	289
Automatic Disk-Based Backup and Recovery (Flash Recovery Area) .....	294
OS Facilities: Online Backup .....	295
Export and Data Pump .....	297
Commercially Available Backup Utilities .....	300
Backup Strategies .....	300
When and What to Back Up .....	301
Examples of Backup Strategies .....	303
Summary .....	304
<b>12 Restore .....</b>	<b>305</b>
Recovering the Database .....	306
Recovery vs. Restoration .....	306
How to Recover Your Database .....	307
Instance Recovery .....	307
Media Recovery .....	308
Log File Recovery .....	309
Point-in-Time Recovery .....	310
Using the Flashback Features of Oracle 10g .....	311
Developing a Recovery Strategy .....	318
Planning for Instance Recovery .....	318
Planning for Media Recovery .....	319
Planning for Operator Error Recovery .....	320
Summary .....	320
<b>13 Oracle Data Guard on Linux .....</b>	<b>323</b>
Data Guard Concepts .....	324
Data Guard Services .....	326
Physical Standby Database .....	329
Logical Standby Database .....	330
Configuring and Managing Oracle Data Guard .....	332
Physical Standby Database Configuration .....	332
Logical Standby Database Configuration .....	338
Summary .....	340

<b>14 Oracle Advanced Replication on Linux .....</b>	<b>341</b>
Introduction to Oracle Replication .....	342
Multimaster Replication .....	343
Materialized View Replication .....	345
Hybrid Replication .....	345
Oracle Replication on Linux .....	346
I/O Capacity .....	346
Database Layout .....	346
Sufficient Network Capacity .....	349
Summary .....	349
<b>15 Configuring the Hardware for Linux and Oracle .....</b>	<b>351</b>
Hardware Tuning Fundamentals .....	352
Choosing the Right Hardware .....	352
Tuning Hardware for Linux .....	359
Network Tuning on Linux .....	365
Choosing the Right Network .....	366
Tuning Linux for Network Performance .....	366
Summary .....	367
<b>16 Tuning SQL .....</b>	<b>369</b>
Examining the Execution Plan .....	370
Plan Table Review .....	370
Viewing the Execution Plan .....	374
Understanding the Explain Plan .....	377
Understanding the Optimizer Decision Process .....	380
A Working Example .....	382
Controlling the Optimizer Decisions .....	385
SQL Tuning Tips .....	388
Summary .....	401
<b>17 Indexes and Partitioning .....</b>	<b>403</b>
Introduction to Oracle 10g Indexing Options .....	404
Oracle 10g Index Options .....	406
The Index Clustering Factor .....	415
Oracle 10g Partitioning .....	417
Index Partitioning Options .....	421
IOT Partitioning Options .....	423
Using Partitioned Indexes for Performance .....	423
Summary .....	424
PART V	
<b>Oracle Products on Linux</b>	
<b>18 Oracle Application Server 10g on Linux .....</b>	<b>427</b>
What Components Make Up Oracle Application Server 10g? .....	428
Java 2 Enterprise Edition .....	429

Management and Security .....	429
Portals .....	430
Wireless .....	430
Business Intelligence .....	430
E-Business Integration .....	430
Using Oracle Application Server 10g .....	431
Installing and Configuring Oracle Application Server 10g on Linux .....	438
Installation .....	438
Configuration .....	438
Summary .....	442
<b>19 The Oracle E-Business Suite on Linux .....</b>	<b>443</b>
Overview of the Oracle E-Business Suite on Linux .....	444
What Is the Oracle E-Business Suite? .....	445
E-Business Suite Architecture .....	445
What Modules Are Available on Linux .....	448
Using the Oracle E-Business Suite on Linux .....	448
The E-Business Suite in a RAC Environment .....	448
Installing and Configuring the E-Business Suite on Linux .....	449
Installation .....	449
Migrating to the E-Business Suite on Linux .....	451
Migrating the Application Tier .....	452
Migrating the Database Tier .....	453
Converting the Database Tier to RAC .....	454
Testing .....	455
Summary .....	455
<b>Index .....</b>	<b>457</b>