1. Record Nr. UNINA9910956707003321

Autore	Fiorillo Ciro
Titolo	Oracle Database 11gR2 performance tuning cookbook : over 80 recipes to help beginners achieve better performance from Oracle Database applications / / Ciro Fiorillo
Pubbl/distr/stampa	Birmingham, U.K., : Packt Enterprise, 2012
ISBN	9786613453440 9781283453448 1283453444 9781849682619 1849682615
Edizione	[1st edition]
Descrizione fisica	1 online resource (543 p.)
Disciplina	005.75/85 005.7575
Soggetti	Relational databases Database management
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index. "Quick answers to common problems"--Cover.
Nota di bibliografia	Includes index.
Nota di contenuto	Cover; Copyright; Credits; About the Author; Acknowledgement; About the Reviewers; www.PacktPub.com; Table of Contents; Preface; Chapter 1: Starting with Performance Tuning; Introduction; Reviewing the performance tuning process; Exploring the example database; Acquiring data using a data dictionary and dynamic performance views; Analyzing data using Statspack reports; Diagnosing performance issues using the alert log; Analyzing data using Automatic Workload Repository (AWR); Analyzing data using Automatic Database Diagnostic Monitor (ADDM); A working example Chapter 2: Optimizing Application DesignIntroduction; Optimizing connection management; Improving performance sharing reusable code; Reducing the number of requests to the database using stored procedures; Reducing the number of requests to the database using sequences; Reducing the number of requests to the database using materialized views; Optimizing performance with schema

denormalization; Avoiding dynamic SQL; Chapter 3: Optimizing Storage Structures; Introduction; Avoiding row chaining; Avoiding row migration; Using LOBs; Using index clusters; Using hash clusters
Indexing the correct wayRebuilding index; Compressing indexes; Using reverse key indexes; Using bitmap indexes; Migrating to index organized tables; Using partitioning; Chapter 4: Optimizing SQL Code; Introduction; Using bind variables; Avoiding full table scans; Exploring index lookup; Exploring index skip-scan and index range-scan; Introducing arrays and bulk operations; Optimizing joins; Using subqueries; Tracing SQL activity with SQL Trace and TKPROF; Chapter 5: Optimizing Sort Operations; Introduction; Sorting-in-memory and on-disk; Sorting and indexing; Writing top n queries and ranking
Using count, min/max, and group-byAvoiding sorting in set operations: union, minus, and intersect; Troubleshooting temporary tablespaces; Chapter 6: Optimizing PL/SQL Code; Introduction; Using bind variables and parsing; Array processing and bulk-collect; Passing values with NOCOPY (or not); Using short-circuit IF statements; Avoiding recursion; Using native compilation; Taking advantage of function result cache; Inlining PL/SQL code; Using triggers and virtual columns; Chapter 7: Improving the Oracle Optimizer; Introduction; Exploring optimizer hints; Collecting statistics; Using histograms
Managing stored outlinesIntroducing Adaptive Cursor Sharing for bind variable peeking; Creating SQL Tuning Sets; Using the SQL Tuning Advisor; Configuring and using SQL Baselines; Chapter 8: Other Optimizations; Introduction; Caching results with the client-side result cache; Enabling parallel SQL; Direct path inserting; Using create table as select; Inspecting indexes and triggers overhead; Loading data with SQL*Loader and Data Pump; Chapter 9: Tuning Memory; Introduction; Tuning memory to avoid Operating System paging; Tuning the Library Cache; Tuning the Shared Pool
Tuning the Program Global Area and the User Global Area

Sommario/riassunto

In this book you will find both examples and theoretical concepts covered. Every recipe is based on a script/procedure explained step-by-step, with screenshots, while theoretical concepts are explained in the context of the recipe, to explain why a solution performs better than another. This book is aimed at software developers, software and data architects, and DBAs who are using or are planning to use the Oracle Database, who have some experience and want to solve performance problems faster and in a rigorous way. If you are an architect who wants to design better applications, a DBA who is