


Relational Database Index Design and the Optimizers

Von Tapio Lahdenmaki, Mike Leach

ePub | *DOC | audiobook | ebooks | Download PDF



DOWNLOAD



READ ONLINE

Produktinformation -Verkaufsrang: #1403057 in eBooksVerffentlicht am: 2007-08-10Erscheinungsdatum: 2007-08-10File Name: B000PY48KE | File size: 41.Mb

Von Tapio Lahdenmaki, Mike Leach : Relational Database Index Design and the Optimizers before purchasing it in order to gage whether or not it would be worth my time, and all praised Relational Database Index Design and the Optimizers:

KundenrezensionenHilfreichste Kundenrezensionen0 von 2 Kunden fanden die folgende Rezension hilfreich. Ein MUSS fr performance and index bebesiterteVon Jiri SaryAlle, die sich in den Index-Design und die damit

verbundene Performance auf relationalen Datenbanken einarbeiten und vertiefen möchten, ist dieses Buch absolut unerlässlich! Strategien und Konzepte werden gut und genau erklärt. Auch diejenigen, welche bereits ein Grundwissen in diesem Bereich haben, bekommen hier viele und sehr nützliche Informationen. Und das erst noch so ausgelegt, dass es nicht von Bedeutung ist, ob die Datenbank DB2, Oracle, SQL Server oder eine andere ist.

Kurzbeschreibung Improve the performance of relational databases with indexes designed for today's hardware Over the last few years, hardware and software have advanced beyond all recognition, so it's hardly surprising that relational database performance now receives much less attention. Unfortunately, the reality is that the improved hardware hasn't kept pace with the ever-increasing quantity of data processed today. Although disk packing densities have increased enormously, making storage costs extremely low and sequential read very fast, random reads are still painfully slow. Many of the old design recommendations are therefore no longer valid-the optimal point of indexing has come a long way. Consequently many of the old problems haven't actually gone away-they have simply changed their appearance. This book provides an easy but effective approach to the design of indexes and tables. Using lots of examples and case studies, the authors describe how the DB2, Oracle, and SQL Server optimizers determine how to access data, and how CPU and response times for the resulting access paths can be quickly estimated. This enables comparisons to be made of the various designs, and helps you choose available choices for the most appropriate design. This book is intended for anyone who wants to understand the issues of SQL performance or how to design tables and indexes effectively. With this title, readers with many years of experience of relational systems will be able to better grasp the implications that have been brought into play by the introduction of new hardware. **Pressestimmen** "I recommend this book to all those who have anything to do with database performance. It is a must-read for all database administrations, database designers, performance-tuning specialists, and application programmers..." (Computing s.com, November 20, 2005) "I recommend this book to all those who have anything to do with database performance. It is a must-read for all database administrations, database designers, performance-tuning specialists, and application programmers..." (Computing s.com, November 20, 2005) **Kurzbeschreibung** Improve the performance of relational databases with indexes designed for today's hardware Over the last few years, hardware and software have advanced beyond all recognition, so it's hardly surprising that relational database performance now receives much less attention. Unfortunately, the reality is that the improved hardware hasn't kept pace with the ever-increasing quantity of data processed today. Although disk packing densities have increased enormously, making storage costs extremely low and sequential read very fast, random reads are still painfully slow. Many of the old design recommendations are therefore no longer valid-the optimal point of indexing has come a long way. Consequently many of the old problems haven't actually gone away-they have simply changed their appearance. This book provides an easy but effective approach to the design of indexes and tables. Using lots of examples and case studies, the authors describe how the DB2, Oracle, and SQL Server optimizers determine how to access data, and how CPU and response times for the resulting access paths can be quickly estimated. This enables comparisons to be made of the various designs, and helps you choose available choices for the most appropriate design. This book is intended for anyone who wants to understand the issues of SQL performance or how to design tables and indexes effectively. With this title, readers with many years of experience of relational systems will be able to better grasp the implications that have been brought into play by the introduction of new hardware.