
Read Online SQL Performance Explained

Recognizing the habit ways to get this books **SQL Performance Explained** is additionally useful. You have remained in right site to begin getting this info. acquire the SQL Performance Explained join that we give here and check out the link.

You could buy guide SQL Performance Explained or acquire it as soon as feasible. You could quickly download this SQL Performance Explained after getting deal. So, behind you require the book swiftly, you can straight get it. Its hence agreed easy and for that reason fats, isnt it? You have to favor to in this appearance

O98L9P - LOGAN DAISY

Troubleshoot query performance issues, identify anti-patterns in code, and write efficient T-SQL queries Key Features Discover T-SQL functionalities and services that help you interact with relational database- Understand the roles, tasks and responsibilities of a T-SQL developer Explore solutions for carrying out database querying tasks, database administration, and troubleshooting Book Description Transact-SQL (T-SQL) is Microsoft's proprietary extension to the SQL language that is used with Microsoft SQL Server and Azure SQL Database. This book will be a useful guide to learning the art of writing efficient T-SQL code in modern SQL Server versions, as well as the Azure SQL Database. The

book will get you started with query processing fundamentals to help you write powerful, performant T-SQL queries. You will then focus on query execution plans and learn how to leverage them for troubleshooting. In the later chapters, you will learn how to identify various T-SQL patterns and anti-patterns. This will help you analyze execution plans to gain insights into current performance, and determine whether or not a query is scalable. You will also learn to build diagnostic queries using dynamic management views (DMVs) and dynamic management functions (DMFs) to address various challenges in T-SQL execution. Next, you will study how to leverage the built-in tools of SQL Server to shorten the time taken to address query performance and scalability issues. In the con-

cluding chapters, the book will guide you through implementing various features, such as Extended Events, Query Store, and Query Tuning Assistant using hands-on examples. By the end of this book, you will have the skills to determine query performance bottlenecks, avoid pitfalls, and discover the anti-patterns in use. Foreword by Conor Cunningham, Partner Architect - SQL Server and Azure SQL - Microsoft What you will learn Use Query Store to understand and easily change query performance Recognize and eliminate bottlenecks that lead to slow performance Deploy quick fixes and long-term solutions to improve query performance Implement best practices to minimize performance risk using T-SQL Achieve optimal performance by ensuring careful query and in-

dex design Use the latest performance optimization features in SQL Server 2017 and SQL Server 2019 Protect query performance during upgrades to newer versions of SQL Server Who this book is for This book is for database administrators, database developers, data analysts, data scientists, and T-SQL practitioners who want to get started with writing T-SQL code and troubleshooting query performance issues, through the help of practical examples. Previous knowledge of T-SQL querying is not required to get started on this book.

SQL is full of difficulties and traps for the unwary. You can avoid them if you understand relational theory, but only if you know how to put the theory into practice. In this insightful book, author C.J. Date explains relational theory in depth, and demonstrates through numerous examples and exercises how you can apply it directly to your use of SQL. This second edition includes new material on recursive queries, “missing information” without nulls, new update operators, and topics such as aggregate operators, grouping and ungrouping, and view updating. If you have a modest-to-advanced background in

SQL, you’ll learn how to deal with a host of common SQL dilemmas. Why is proper column naming so important? Nulls in your database are causing you to get wrong answers. Why? What can you do about it? Is it possible to write an SQL query to find employees who have never been in the same department for more than six months at a time? SQL supports “quantified comparisons,” but they’re better avoided. Why? How do you avoid them? Constraints are crucially important, but most SQL products don’t support them properly. What can you do to resolve this situation? Database theory and practice have evolved since the relational model was developed more than 40 years ago. SQL and Relational Theory draws on decades of research to present the most up-to-date treatment of SQL available. C.J. Date has a stature that is unique within the database industry. A prolific writer well known for the bestselling textbook *An Introduction to Database Systems* (Addison-Wesley), he has an exceptionally clear style when writing about complex principles and theory.

A very practical guide to making databases run faster and better. A poorly perform-

ing database application can cost each user time, and have an impact on other applications running on the same computer or the same network. This book will help DBAs and programmers improve the performance of their databases.

The goal of Defensive Programming is to produce resilient code that responds gracefully to the unexpected. To the SQL Server programmer, this means T-SQL code that behaves consistently and predictably in cases of unexpected usage, doesn’t break under concurrent loads, and survives predictable changes to database schemas and settings. Inside this book, you will find dozens of practical, defensive programming techniques that will improve the quality of your T-SQL code and increase its resilience and robustness.

Bert Scalzo and Dan Hotka have written the definitive, up-to-date guide to Version 12.x, Dell’s powerful new release of Toad for Oracle. Packed with step-by-step recipes, detailed screen shots, and hands-on exercises, *Toad for Oracle Unleashed* shows both developers and DBAs how to maximize their productivity. Drawing on their unsurpassed experience running

Toad in production Oracle environments, Scalzo and Hotka thoroughly cover every area of Toad's functionality. You'll find practical insights into each of Toad's most useful tools, from App Designer to Doc Generator, ER Diagrammer to Code Road Map. The authors offer proven solutions you can apply immediately to solve a wide variety of problems, from maintaining code integrity to automating performance and scalability testing. Learn how to... Install and launch Toad, connect to a database, and explore Toad's new features Customize Toad to optimize productivity in your environment Use the Editor Window to execute SQL and PL/SQL, and view, save, or convert data Browse your schema, and create and edit objects Quickly generate useful reports with FastReport and Report Manager Clarify your database's tables and data with the powerful Entity Relationship Diagrammer (ERD) and HTML documentation generator Work more efficiently with PL/SQL using code templates, snippets, and shortcuts Automate actions and applications with Automation Designer Perform key DBA tasks including database health checks, tablespace management, database and schema

comparisons, and object rebuilding Identify and optimize poorlyperforming SQL and applications ON THE WEB: Download all examples and source code presented in this book from informit.com/title/9780134131856 as it becomes available.

Provides information on extracting Oracle data to create and format reports.

How can you bring out MySQL's full power? With High Performance MySQL, you'll learn advanced techniques for everything from designing schemas, indexes, and queries to tuning your MySQL server, operating system, and hardware to their fullest potential. This guide also teaches you safe and practical ways to scale applications through replication, load balancing, high availability, and failover. Updated to reflect recent advances in MySQL and InnoDB performance, features, and tools, this third edition not only offers specific examples of how MySQL works, it also teaches you why this system works a.

Design and configure SQL Server instances and databases in support of high-throughput applications that are mission-critical and provide consistent response times in the face of variations in user numbers and

query volumes. Learn to configure SQL Server and design your databases to support a given instance and workload. You'll learn advanced configuration options, in-memory technologies, storage and disk configuration, and more, all toward enabling your desired application performance and throughput. Configuration doesn't stop with implementation. Workloads change over time, and other impediments can arise to thwart desired performance. High Performance SQL Server covers monitoring and troubleshooting to aid in detecting and fixing production performance problems and minimizing application outages. You'll learn a variety of tools, ranging from the traditional wait analysis methodology to the new query store, and you'll learn how improving performance is really an iterative process. High Performance SQL Server is based on SQL Server 2016, although most of its content can be applied to prior versions of the product. This book is an excellent complement to performance tuning books focusing on SQL queries, and provides the other half of what you need to know by focusing on configuring the instances on which mission-critical queries are executed. Covers

SQL Server instance-configuration for optimal performance Helps in implementing SQL Server in-memory technologies Provides guidance toward monitoring and ongoing diagnostics What You Will Learn Understand SQL Server's database engine and how it processes queries Configure instances in support of high-throughput applications Provide consistent response times to varying user numbers and query volumes Design databases for high-throughput applications with focus on performance Record performance baselines and monitor SQL Server instances against them Troubleshoot and fix performance problems Who This Book Is For SQL Server database administrators, developers, and data architects. The book is also of use to system administrators who are managing and are responsible for the physical servers on which SQL Server instances are run.

T-SQL insiders help you tackle your toughest queries and query-tuning problems Squeeze maximum performance and efficiency from every T-SQL query you write or tune. Four leading experts take an in-depth look at T-SQL's internal architec-

ture and offer advanced practical techniques for optimizing response time and resource usage. Emphasizing a correct understanding of the language and its foundations, the authors present unique solutions they have spent years developing and refining. All code and techniques are fully updated to reflect new T-SQL enhancements in Microsoft SQL Server 2014 and SQL Server 2012. Write faster, more efficient T-SQL code: Move from procedural programming to the language of sets and logic Master an efficient top-down tuning methodology Assess algorithmic complexity to predict performance Compare data aggregation techniques, including new grouping sets Efficiently perform data-analysis calculations Make the most of T-SQL's optimized bulk import tools Avoid date/time pitfalls that lead to buggy, poorly performing code Create optimized BI statistical queries without additional software Use programmable objects to accelerate queries Unlock major performance improvements with In-Memory OLTP Master useful and elegant approaches to manipulating graphs About This Book For experienced T-SQL practitioners Includes coverage updated from Inside Microsoft SQL

Server 2008 T-SQL Querying and Inside Microsoft SQL Server 2008 T-SQL Programming Valuable to developers, DBAs, BI professionals, and data scientists Covers many MCSE 70-464 and MCSA/MCSE 70-461 exam topics

Queries not running fast enough? Wondering about the in-memory database features in 2014? Tired of phone calls from frustrated users? Grant Fritchey's book SQL Server Query Performance Tuning is the answer to your SQL Server query performance problems. The book is revised to cover the very latest in performance optimization features and techniques, especially including the newly-added, in-memory database features formerly known under the code name Project Hekaton. This book provides the tools you need to approach your queries with performance in mind. SQL Server Query Performance Tuning leads you through understanding the causes of poor performance, how to identify them, and how to fix them. You'll learn to be proactive in establishing performance baselines using tools like Performance Monitor and Extended Events. You'll learn to recognize bottlenecks and defuse them before the phone rings. You'll learn some

quick solutions too, but emphasis is on designing for performance and getting it right, and upon heading off trouble before it occurs. Delight your users. Silence that ringing phone. Put the principles and lessons from SQL Server Query Performance Tuning into practice today. Covers the in-memory features from Project Hekaton Helps establish performance baselines and monitor against them Guides in troubleshooting and eliminating of bottlenecks that frustrate users

Take a deep dive into perhaps the single most important facet of good performance: indexes, and how to best use them. Recent updates to SQL Server have made it possible to create indexes in situations that in the past would have prevented their use. Other improvements covered in this book include new dynamic management views, the ability to pause and resume index maintenance, and the ability to more easily recover from failures during index creation and maintenance operations. This new edition also brings new content around the indexing of columnstore and in-memory tables, showing how these new types of tables and the queries that execute against them can also benefit

from good indexing practices. The book begins with explanations of the types of indexes and how they are stored in databases. Moving deeper into the topic, and further into the book, you will look at the statistics that are accumulated both by indexes and on indexes. You will better understand what indexes are doing in the database and what can be done to mitigate and improve their effect on performance. You will get a look at the Index Advisor now available in Azure SQL Database, and learn how to review and maintain the health of your indexes. The final chapters present a guided tour through a number of scenarios showing approaches you can take to investigate, mitigate, and improve the performance of your database. What You Will Learn Properly index row store, columnstore, and in-memory tables Review statistics to understand indexing choices made by the optimizer Apply indexing strategies such as covering indexes, included columns, and index intersections Recognize and remove unnecessary indexes Design effective indexes for full-text, spatial, and XML data types Manage the big picture: Encompass all indexes in a database, and all database ins-

tances on a server Who This Book Is For Database administrators and developers who are ready to lift the performance of their database environment by thoughtfully building indexes to speed up queries that matter the most and make a difference to the business

Arguably the most capable of all the open source databases, PostgreSQL is an object-relational database management system first developed in 1977 by the University of California at Berkeley. In spite of its long history, this robust database suffers from a lack of easy-to-use documentation. Practical PostgreSQL fills that void with a fast-paced guide to installation, configuration, and usage. This comprehensive new volume shows you how to compile PostgreSQL from source, create a database, and configure PostgreSQL to accept client-server connections. It also covers the many advanced features, such as transactions, versioning, replication, and referential integrity that enable developers and DBAs to use PostgreSQL for serious business applications. The thorough introduction to PostgreSQL's PL/pgSQL programming language explains how you can use this very useful but under-documented fea-

ture to develop stored procedures and triggers. The book includes a complete command reference, and database administrators will appreciate the chapters on user management, database maintenance, and backup & recovery. With Practical PostgreSQL, you will discover quickly why this open source database is such a great open source alternative to proprietary products from Oracle, IBM, and Microsoft.

This book describes, diagnoses, and solves the most common problems with SQL Server 2005, 2008, and 2008 R2. The authors explain a basic approach to troubleshooting and the essential tools. They explore areas in which problems arise with regularity: high disk I/O (RAID misconfiguration, inadequate I/O throughput, poor workload distribution, SAN issues, disk partition misalignment); high CPU usage (insufficient memory, poorly written queries, inadequate indexing, inappropriate configuration option settings); memory mismanagement; missing indexes; blocking (caused mainly by poorly designed databases that lack proper keys and indexing, and applications that apply needlessly restrictive transaction isolation levels); deadlocking (Book-

mark Lookup, Serializable Range Scan, Cascading Constraint); full transaction logs (lack of log backups, hefty index maintenance operations, long running transaction, problems with replication and mirroring environments); and accidentally-lost data. Finally, the authors discuss diagnosing tools such as the Performance Monitor, Dynamic Management Views, and server-side tracing. --

Expert Indexing in Oracle Database 11g is about the one database structure at the heart of almost all performance concerns: the index. Database system performance is one of the top concerns in information technology today. Administrators struggle to keep up with the explosion of access and activity driven by the proliferation of computing into everything from phones to tablets to PCs in our increasingly connected world. At the heart of any good-performing database lies a sound indexing strategy that makes appropriate use of indexing, and especially of the vendor-specific indexing features on offer. Few databases fully exploit the wealth of data access mechanisms provided by Oracle. Expert Indexing in Oracle Database 11g helps by bringing together information indexing and how

to use it into one, convenient and blissfully short volume that you can read quickly and have at your fingertips for reference. Learn the different types of indices available and when each is best applied. Recognize when queries aren't using indices as you intend. Manage your indexing for maximum performance. Let Expert Indexing in Oracle Database 11g be your guide to deep mastery of the most fundamental performance optimization structure in Oracle Database. Explains how indices work, how they help, and how they hinder Demystifies the various index choices Describes the database administration chores associated with indices

For all the buzz about trendy IT techniques, data processing is still at the core of our systems, especially now that enterprises all over the world are confronted with exploding volumes of data. Database performance has become a major headache, and most IT departments believe that developers should provide simple SQL code to solve immediate problems and let DBAs tune any bad SQL later. In The Art of SQL, author and SQL expert Stephane Faroult argues that this safe approach only leads to disaster. His insightful book,

named after Art of War by Sun Tzu, contends that writing quick inefficient code is sweeping the dirt under the rug. SQL code may run for 5 to 10 years, surviving several major releases of the database management system and on several generations of hardware. The code must be fast and sound from the start, and that requires a firm understanding of SQL and relational theory. The Art of SQL offers best practices that teach experienced SQL users to focus on strategy rather than specifics. Faroult's approach takes a page from Sun Tzu's classic treatise by viewing database design as a military campaign. You need knowledge, skills, and talent. Talent can't be taught, but every strategist from Sun Tzu to modern-day generals believed that it can be nurtured through the experience of others. They passed on their experience acquired in the field through basic principles that served as guiding stars amid the sound and fury of battle. This is what Faroult does with SQL. Like a successful battle plan, good architectural choices are based on contingencies. What if the volume of this or that table increases unexpectedly? What if, following a merger, the number of users doubles? What if you want to keep

several years of data online? Faroult's way of looking at SQL performance may be unconventional and unique, but he's deadly serious about writing good SQL and using SQL well. The Art of SQL is not a cookbook, listing problems and giving recipes. The aim is to get you-and your manager-to raise good questions.

If a query is performing poorly, and you can't understand why, then that query's execution plan will tell you not only what data set is coming back, but also what SQL Server did, and in what order, to get that data. It will reveal how the data was retrieved, and from which tables and indexes, what types of joins were used, at what point filtering, sorting and aggregation occurred, and a whole lot more. These details will often highlight the likely source of any problem. I wrote this book with the singular goal of teaching you how to read SQL Server Execution plans It will explain, among many other things, the following: How to capture execution plans using manual and automatic methods A documented method for reading and interpreting execution plans How common SQL Server objects, such as indexes, views, stored procedures, and so on, appear in execution

plans How to control execution plans with hints and plan guides, and why this is a double-edged sword How the Query Store works with, and collects data on, execution plans With this knowledge, you'll have everything you need to read the execution plan, for any query of your own, regardless of complexity, and understand what it does and what is causing the bad performance. It is still your job to work out how best to fix it, but your new understanding of execution plans will give a much better chance of success!

A guide to SQL covers such topics as retrieving records, metadata queries, working with strings, data arithmetic, date manipulation, reporting and warehousing, and hierarchical queries.

Mastering PostgreSQL in Application Development is intended for developers working on applications that use a database server. The book addresses specifically the PostgreSQL RDBMS: it actually is the world's most advanced Open Source database as said in its slogan on the official website. By the end of this book, you will know why, and agree!

A poorly performing database application

not only costs users time, but also has an impact on other applications running on the same computer or the same network. SQL Tuning provides an essential next step for SQL developers and database administrators who want to extend their SQL tuning expertise and get the most from their database applications. There are two basic issues to focus on when tuning SQL: how to find and interpret the execution plan of an SQL statement and how to change SQL to get a specific alternate execution plan. SQL Tuning provides answers to these questions and addresses a third issue that's even more important: how to find the optimal execution plan for the query to use. Author Dan Tow outlines a timesaving method he's developed for finding the optimum execution plan--rapidly and systematically--regardless of the complexity of the SQL or the database platform being used. You'll learn how to understand and control SQL execution plans and how to diagram SQL queries to deduce the best execution plan for a query. Key chapters in the book include exercises to reinforce the concepts you've learned. SQL Tuning concludes by addressing special concerns and unique solutions to "unsolv-

able problems." Whether you are a programmer who develops SQL-based applications or a database administrator or other who troubleshoots poorly tuned applications, SQL Tuning will arm you with a reliable and deterministic method for tuning your SQL queries to gain optimal performance.

When your database application isn't running fast enough, troubleshooting is usually your first move. Finding the slow part of an application is often easy, but discovering a solution can prove much more difficult. Troubleshooting Oracle Performance helps by providing a systematic approach to addressing the underlying causes of poor database application performance. Written for developers by an application developer who has learned by doing, this book shows you how to plan for performance as you would for any other application requirement.

Execution plans show you what's going on behind the scenes in SQL Server. They can provide you with a wealth of information on how your queries are being executed by SQL Server, including: Which indexes are being used, and where no indexes are

being used at all. How the data is being retrieved, and joined, from the tables defined in your query. How aggregations in GROUP BY queries are put together. The anticipated load and the estimated cost that all these operations place upon the system. Grant Fritchey's book is the only in-depth look at how to improve your SQL query performance through careful design of execution plans. Sample chapters of the ebook have garnered stunning reviews, such as: "All I can say is WOW. This has to be the best reference I have ever seen on Execution Plans in SQL Server. My hats off to Grant Fritchey" Jonathan Kehayias.

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.

Get the most out of the rich development capabilities of SQL Server 2016 to build efficient database applications for your organization About This Book Utilize the new

enhancements in Transact-SQL and security features in SQL Server 2016 to build efficient database applications Work with temporal tables to get information about data stored in the table at any point in time A detailed guide to SQL Server 2016, introducing you to multiple new features and enhancements to improve your overall development experience Who This Book Is For This book is for database developers and solution architects who plan to use the new SQL Server 2016 features for developing efficient database applications. It is also ideal for experienced SQL Server developers who want to switch to SQL Server 2016 for its rich development capabilities. Some understanding of the basic database concepts and Transact-SQL language is assumed. What You Will Learn Explore the new development features introduced in SQL Server 2016 Identify opportunities for In-Memory OLTP technology, significantly enhanced in SQL Server 2016 Use columnstore indexes to get significant storage and performance improvements Extend database design solutions using temporal tables Exchange JSON data between applications and SQL Server in a more efficient way Migrate historical data transparently

and securely to Microsoft Azure by using Stretch Database Use the new security features to encrypt or to have more granular control over access to rows in a table Simplify performance troubleshooting with Query Store Discover the potential of R's integration with SQL Server In Detail Microsoft SQL Server 2016 is considered the biggest leap in the data platform history of the Microsoft, in the ongoing era of Big Data and data science. Compared to its predecessors, SQL Server 2016 offers developers a unique opportunity to leverage the advanced features and build applications that are robust, scalable, and easy to administer. This book introduces you to new features of SQL Server 2016 which will open a completely new set of possibilities for you as a developer. It prepares you for the more advanced topics by starting with a quick introduction to SQL Server 2016's new features and a recapitulation of the possibilities you may have already explored with previous versions of SQL Server. The next part introduces you to small delights in the Transact-SQL language and then switches to a completely new technology inside SQL Server - JSON support. We also take a look at the Stretch database,

security enhancements, and temporal tables. The last chapters concentrate on implementing advanced topics, including Query Store, columnstore indexes, and In-Memory OLTP. You will finally be introduced to R and how to use the R language with Transact-SQL for data exploration and analysis. By the end of this book, you will have the required information to design efficient, high-performance database applications without any hassle. Style and approach This book is a detailed guide to mastering the development features offered by SQL Server 2016, with a unique learn-as-you-do approach. All the concepts are explained in a very easy-to-understand manner and are supplemented with examples to ensure that you—the developer—are able to take that next step in building more powerful, robust applications for your organization with ease.

Written by a Senior Database Administrator who has worked with the Oracle RDBMS for thirty years, this is a book which teaches the skill of SQL Tuning for the Oracle Database. Not a list of one-off tricks or tips, nor a glossing over of topics; this book offers an in-depth process covering discovery, analysis, and problem reso-

lution. Learn the science behind SQL Tuning. Learn and apply the FILTERED ROWS PERCENTAGE Cardinality based method of tuning Determine a query's Driving Table and Join Order Construct Query Diagrams, Data Models, and Join Trees Build and use Count / Filter / and Reconstruction Queries Identify Waste in a Query Execution Plan Zero in on Cardinality Divergence using Estimated vs. Actuals Use the ACCESS / FILTER / COVERAGE strategy to build indexes for Problem Queries Exploit THE 2% RULE in analyzing Access method and Join method Classify queries as Precision Style or Warehouse Style Understand Hash Join mechanics and make Hash Joins go faster Make HINTS work as Detection Tools rather than clubs Avoid early Database Design flaws Manage Statistics and deal with common Statistics problems (NDV, Uniform Distribution, Independence, Dynamic Sampling) (Staleness, Skew, Dependence, Defaulting, Out-Of-Bounds, Transiency, Bloat) Perfect your Question Based Analysis Technique and more Included are: a special chapter for EXADATA, a LAB which demonstrates the cardinality based process of SQL Tuning, and twenty three magical SQL scripts that make the process of SQL

Tuning easy to do. Learn the skill of SQL Tuning as taught by an expert who does it for a living, and become the go-to specialist in your company. Chapter 1: DRIVING TABLE and JOIN ORDER Chapter 2: Ways to Use a Query Execution Plan Chapter 3: The Best Indexes for a Query Chapter 4: JOINS Chapter 5: HINTS Chapter 6: BASICS Chapter 7: ROW COUNTS and RUN TIMES Chapter 8: EXADATA LAB: Reverse Engineering the QEP Appendix: Know Your Scripts Scripts for analyzing queries and plans Scripts for examining an active database Scripts for looking at metadata showplan showplanshort showplanconstraints showplancountqueries showplandatatamodel showplandrivingtable showplanfilterqueries showplanfrpspreadsheetcode showplanindexes showplannumrows showplanquerydiagram showplantables showplantablesunique loadplanfromcache loadplanfromhist showtopcpu showowner showindexes showconstraints showcolstats showhistograms showallscanrates showallworkareas It's all about the Cardinalities

This is the second edition of the popular practitioner's guide to SQL, the industry-s-

standard database query language. Like most computer languages, SQL can be overwhelming when you first see it, but for years readers have relied on this book to clear the confusion and explain how SQL works and how to use it effectively. Packed with tips, tricks, and good information, *SQL Clearly Explained, Second Edition* teaches database users and programmers everything they need to know to get their job done including

- formulating SQL queries,
- understanding how queries are processed by the DBMS,
- maximizing performance,
- using SQL to enter, modify, or delete data,
- creating and maintaining database structural elements, and
- embedding SQL in applications.

Features

- Updated and expanded to include changes in the SQL standard (SQL:1999) as well as recently implemented aspects of SQL-92.
- Includes CD with examples from the book as well as MySQL, a popular open-source DBMS, on which the examples are based.
- Web enhanced with extra features available online at www.mkp.com.

* Second edition of classic SQL handbook
* Updated to cover changes in the SQL language standard (SQL:1999)
* Includes CD with MySQL software

How can you realize MySQL's full power? With *High Performance MySQL*, you'll learn advanced techniques for everything from setting service-level objectives to designing schemas, indexes, and queries to tuning your server, operating system, and hardware to achieve your platform's full potential. This guide also teaches database administrators safe and practical ways to scale applications through replication, load balancing, high availability, and failover. Updated to reflect recent advances in cloud- and self-hosted MySQL, InnoDB performance, and new features and tools, this revised edition helps you design a relational data platform that will scale with your business. You'll learn best practices for database security along with hard-earned lessons in both performance and database stability. Dive into MySQL's architecture, including key facts about its storage engines. Learn how server configuration works with your hardware and deployment choices. Make query performance part of your software delivery process. Examine enhancements to MySQL's replication and high availability. Compare different MySQL offerings in managed cloud environments. Explore MySQL's full stack optimization from

application-side configuration to server tuning. Turn traditional database management tasks into automated processes.

A hands-on resource for SQL Server 2008 troubleshooting methods and tools. SQL Server administrators need to ensure that SQL Server remains running 24/7. Authored by leading SQL Server experts and MVPs, this book provides in-depth coverage of best practices based on a deep understanding of the internals of both SQL Server and the Windows operating system. You'll get a thorough look at the SQL Server database architecture and internals as well as Windows OS internals so that you can approach troubleshooting with a solid grasp of the total processing environment. Armed with this comprehensive understanding, readers will then learn how to use a suite of tools for troubleshooting performance problems whether they originate on the database server or operating system side. Topics Covered: SQL Server Architecture Understanding Memory SQL Server Waits and Extended Events Working with Storage CPU and Query Processing Locking and Latches Knowing Tempdb Defining Your Approach To Troubleshooting Viewing Server Performance with Perf-

Mon and the PAL Tool Tracing SQL Server with SQL Trace and Profiler Consolidating Data Collection with SQLDiag and the Perf-Stats Script Introducing RML Utilities for Stress Testing and Trace File Analysis Bringing It All Together with SQL Nexus Using Management Studio Reports and the Performance Dashboard Using SQL Server Management Data Warehouse Shortcuts to Efficient Data Collection and Quick Analysis Note: CD-ROM/DVD and other supplementary materials are not included as part of eBook file.

The #1 Easy, Common-Sense Guide to SQL Queries—Updated for Today’s Databases, Standards, and Challenges SQL Queries for Mere Mortals® has earned worldwide praise as the clearest, simplest tutorial on writing effective SQL queries. The authors have updated this hands-on classic to reflect new SQL standards and database applications and teach valuable new techniques. Step by step, John L. Viescas and Michael J. Hernandez guide you through creating reliable queries for virtually any modern SQL-based database. They demystify all aspects of SQL query writing, from simple data selection and filtering to joining multiple tables and modifying sets

of data. Three brand-new chapters teach you how to solve a wide range of challenging SQL problems. You’ll learn how to write queries that apply multiple complex conditions on one table, perform sophisticated logical evaluations, and think “outside the box” using unlinked tables. Coverage includes -- Getting started: understanding what relational databases are, and ensuring that your database structures are sound -- SQL basics: using SELECT statements, creating expressions, sorting information with ORDER BY, and filtering data using WHERE -- Summarizing and grouping data with GROUP BY and HAVING clauses - - Drawing data from multiple tables: using INNER JOIN, OUTER JOIN, and UNION operators, and working with subqueries -- Modifying data sets with UPDATE, INSERT, and DELETE statements Advanced queries: complex NOT and AND, conditions, if-then-else using CASE, unlinked tables, driver tables, and more Practice all you want with downloadable sample databases for today’s versions of Microsoft Office Access, Microsoft SQL Server, and the open source MySQL database. Whether you’re a DBA, developer, user, or student, there’s no better way to master SQL. informit.com/aw

forMereMortals.com

Updated for the latest database management systems -- including MySQL 6.0, Oracle 11g, and Microsoft's SQL Server 2008 - - this introductory guide will get you up and running with SQL quickly. Whether you need to write database applications, perform administrative tasks, or generate reports, Learning SQL, Second Edition, will help you easily master all the SQL fundamentals. Each chapter presents a self-contained lesson on a key SQL concept or technique, with numerous illustrations and annotated examples. Exercises at the end of each chapter let you practice the skills you learn. With this book, you will: Move quickly through SQL basics and learn several advanced features Use SQL data statements to generate, manipulate, and retrieve data Create database objects, such as tables, indexes, and constraints, using SQL schema statements Learn how data sets interact with queries, and understand the importance of subqueries Convert and manipulate data with SQL's built-in functions, and use conditional logic in data statements Knowledge of SQL is a must for interacting with data. With Learning SQL,

you'll quickly learn how to put the power and flexibility of this language to work.

Distills key concepts from linear algebra, geometry, matrices, calculus, optimization, probability and statistics that are used in machine learning.

The authors have revised and updated this bestseller to include both the Oracle8i and new Oracle9i Internet-savvy database products.

Learn effective and scalable database design techniques in a SQL Server 2016 and higher environment. This book is revised to cover in-memory online transaction processing, temporal data storage, row-level security, durability enhancements, and other design-related features that are new or changed in SQL Server 2016. Designing an effective and scalable database using SQL Server is a task requiring skills that have been around for forty years coupled with technology that is constantly changing. Pro SQL Server Relational Database Design and Implementation covers everything from design logic that business users will understand, all the way to the physical implementation of design in a SQL Server database. Grounded in best practices and a solid understanding of the underlying

theory, Louis Davidson shows how to "get it right" in SQL Server database design and lay a solid groundwork for the future use of valuable business data. The pace of change in relational database management systems has been tremendous these past few years. Whereas in the past it was enough to think about optimizing data residing on spinning hard drives, today one also must consider solid-state storage as well as data that are constantly held in memory and never written to disk at all except as a backup. Furthermore, there is a trend toward hybrid cloud and on-premise database configurations as well a move toward preconfigured appliances. Pro SQL Server Relational Database Design and Implementation guides in the understanding of these massive changes and in their application toward sound database design. Gives a solid foundation in best practices and relational theory Covers the latest implementation features in SQL Server 2016 Helps you master in-memory OLTP and use it effectively Takes you from conceptual design to an effective, physical implementation What You Will Learn Develop conceptual models of client data using interviews and client documentation Recog-

nize and apply common database design patterns Normalize data models to enhance scalability and the long term use of valuable data Translate conceptual models into high-performing SQL Server databases Secure and protect data integrity as part of meeting regulatory requirements Create effective indexing to speed query performance Who This Book Is For Programmers and database administrators of all types who want to use SQL Server to store data. The book is especially useful to those wanting to learn the very latest design features in SQL Server 2016, features that include an improved approach to in-memory OLTP, durability enhancements, temporal data support, and more. Chapters on fundamental concepts, the language of database modeling, SQL implementation, and of course, the normalization process, lay a solid groundwork for readers who are just entering the field of database design. More advanced chapters serve the seasoned veteran by tackling the very latest in physical implementation features that SQL Server has to offer. The book has been carefully revised to cover all the design-related features that are new in SQL Server 2016.

Explains how to use Structured Query Language to work within a relational database system, including information retrieval, security, data manipulation, and user management.

Pro Oracle SQL unlocks the power of SQL in the Oracle Database—one of the most potent SQL implementations on the market today. To master it requires a three-pronged approach: learn the language features, learn the supporting features that Oracle provides to help use the language effectively, and learn to think and work in sets. Karen Morton and her team help you master powerful aspects of Oracle SQL not found in competing databases. You'll learn analytic functions, the MODEL clause, and advanced grouping syntax—features that will help in creating good queries for reporting and business intelligence applications. Pro Oracle SQL also helps you minimize parsing overhead, read execution plans, test for correct results, and exert control over SQL execution in your database. You'll learn when to create indexes, how to verify that they make a difference, how to use SQL Profiles to optimize SQL in packaged applications, and much more. You'll also understand how

SQL is optimized for working in sets, and that the key to getting accurate results lies in making sure that queries ask clear and precise questions. What's the bottom-line? Pro Oracle SQL helps you work at a truly professional level in Oracle dialect of SQL. You'll master the language, the tools to work effectively with the language, and the right way to think about a problem in SQL. Pro Oracle SQL helps you rise above the crowd to provide stellar service in your chosen profession. Endorsed by the OakTable Network, a group of Oracle technologists well-known for their rigorous and scientific approach to Oracle Database performance Comprehensive—goes beyond the language with a focus on what you need to know to write successful queries and data manipulation statements.

This book provides a comprehensive overview on best practices for troubleshooting and performance tuning in SQL Server. It reviews how to identify performance issues, how to troubleshoot the system in a holistic fashion, and how to properly prioritize tuning efforts in order to induce the best system performance possible. The book also discusses interdependencies between database components,

while spotlighting ways to avoid the bottlenecks that can be triggered by those dependencies. The troubleshooting and performance tuning techniques presented in the book are compatible with any version of SQL Server. They cover both on-premise and Cloud-based SQL Server installations, including Microsoft Azure SQL Databases and Amazon SQL Server RDS. Reflecting the approaches used by many high-end SQL Server consultants, SQL Server Advanced Troubleshooting and Performance Tuning is a valuable resource that will help readers master troubleshooting and performance tuning skills and get the best performance out of SQL Server.

When it comes to choosing, using, and maintaining a database, understanding its internals is essential. But with so many distributed databases and tools available today, it's often difficult to understand what each one offers and how they differ. With this practical guide, Alex Petrov guides developers through the concepts behind modern database and storage engine internals. Throughout the book, you'll explore relevant material gleaned from numerous books, papers, blog posts, and the source code of several open source databases. Th-

ese resources are listed at the end of parts one and two. You'll discover that the most significant distinctions among many modern databases reside in subsystems that determine how storage is organized and how data is distributed. This book examines:

- Storage engines: Explore storage classification and taxonomy, and dive into B-Tree-based and immutable Log Structured storage engines, with differences and use-cases for each
- Storage building blocks: Learn how database files are organized to build efficient storage, using auxiliary data structures such as Page Cache, Buffer Pool and Write-Ahead Log
- Distributed systems: Learn step-by-step how nodes and processes connect and build complex communication patterns
- Database clusters: Which consistency models are commonly used by modern databases and how distributed storage systems achieve consistency

Fully revised, updated, and expanded, *Relational Database Design and Implementation*, Third Edition is the most lucid and effective introduction to the subject available for IT/IS professionals interested in honing their skills in database design, im-

plementation, and administration. This book provides the conceptual and practical information necessary to develop a design and management scheme that ensures data accuracy and user satisfaction while optimizing performance, regardless of experience level or choice of DBMS. The book begins by reviewing basic concepts of databases and database design, then briefly reviews the SQL one would use to create databases. Topics such as the relational data model, normalization, data entities and Codd's Rules (and why they are important) are covered clearly and concisely but without resorting to "Dummies"-style talking down to the reader. Supporting the book's step-by-step instruction are three NEW case studies illustrating database planning, analysis, design, and management practices. In addition to these real-world examples, which include object-relational design techniques, an entirely NEW section consisting of three chapters is devoted to database implementation and management issues. * Principles needed to understand the basis of good relational database design and implementation practices. * Examples to illustrate core concepts for enhanced compre-

hension and to put the book's practical instruction to work. * Methods for tailoring DB design to the environment in which the database will run and the uses to which it will be put. * Design approaches that ensure data accuracy and consistency. * Examples of how design can inhibit or boost database application performance. * Object-relational design techniques, benefits, and examples. * Instructions on how to choose and use a normalization technique. * Guidelines for understanding and applying Codd's rules. * Tools to implement a relational design using SQL. * Techniques for using CASE tools for database design.

Apache Spark is amazing when everything clicks. But if you haven't seen the performance improvements you expected, or still don't feel confident enough to use Spark in production, this practical book is for you. Authors Holden Karau and Rachel Warren demonstrate performance optimizations to help your Spark queries run faster and handle larger data sizes, while using fewer resources. Ideal for software engineers, data engineers, developers, and system administrators working with large-scale data applications, this book describes techniques that can reduce data in-

frastructure costs and developer hours. Not only will you gain a more comprehensive understanding of Spark, you'll also learn how to make it sing. With this book, you'll explore: How Spark SQL's new interfaces improve performance over SQL's

RDD data structure The choice between data joins in Core Spark and Spark SQL Techniques for getting the most out of standard RDD transformations How to work around performance issues in Spark's key/value pair paradigm Writing high-performance Spark code without Scala or the

JVM How to test for functionality and performance when applying suggested improvements Using Spark MLlib and Spark ML machine learning libraries Spark's Streaming components and external community packages