

Mariadb High Performance

A high-performance data access layer must resonate with the underlying database system. Knowing the inner workings of a relational database and the data access frameworks in use can make the difference between a high-performance enterprise application and one that barely crawls. This book is a journey into Java data access performance tuning. From connection management, to batch updates, fetch sizes and concurrency control mechanisms, it unravels the inner workings of the most common Java data access frameworks. The first part aims to reduce the gap between application developers and database administrators. For this reason, it covers both JDBC and the database fundamentals that are of paramount importance when reducing transaction response times. In this first part, you'll learn about connection management, batch updates, statement caching, result set fetching and database transactions. The second part demonstrates how you can take advantage of JPA and Hibernate without compromising application performance. In this second part, you'll learn about the most efficient Hibernate mappings (basic types, associations, inheritance), fetching best practices, caching and concurrency control mechanisms. The third part is dedicated to jOOQ and its powerful type-safe querying capabilities, like window functions, common table expressions, upsert, stored procedures and database functions. This book provides an overview of the resources and research projects that are bringing Big Data and High Performance Computing (HPC) on converging tracks. It demystifies Big Data and HPC for the reader by covering the primary resources, middleware, applications, and tools that enable the usage of HPC platforms for Big Data management and processing. Through interesting use-cases from traditional and non-traditional HPC domains, the book highlights the most critical challenges related to Big Data processing and management, and shows ways to mitigate them using HPC resources. Unlike most books on Big Data, it covers a variety of alternatives to Hadoop, and explains the differences between HPC platforms and Hadoop. Written by professionals and researchers in a range of departments and fields, this book is designed for anyone studying Big Data and its future directions. Those studying HPC will also find the content valuable.

MariaDB is a database that has become very popular in the few short years that it has been around. It does not require a big server or expensive support contract. It is also powerful enough to be the database of choice for some of the biggest and most popular websites in the world, taking full advantage of the latest computing hardware available. From installing and configuring through basic usage and maintenance, each chapter in this revised and expanded guide leads on sequentially and logically from the one before it, introducing topics in their natural order so you learn what you need, when you need it. The book is based on the latest release of MariaDB and covers all the latest features and functions. By the end of this beginner-friendly book, not only will you have a running installation of MariaDB, but you will have practical, hands-on experience in the basics of how to install, configure, administer, use, and maintain it.

Leverage the lethal combination of Docker and Kubernetes to automate deployment and management of Java applications About This Book Master using Docker and Kubernetes to build, deploy and manage Java applications in a jiff Learn how to create your own Docker image and customize your own cluster using Kubernetes Empower the journey from development to production using this practical guide. Who This Book Is For The book is aimed at Java developers who are eager to build, deploy, and manage applications very quickly using container technology. They need have no knowledge of Docker and Kubernetes. What You Will Learn Package Java applications into Docker images Understand the running of containers locally Explore development and deployment options with Docker Integrate Docker into Maven builds Manage and monitor Java applications running on Kubernetes clusters Create Continuous Delivery pipelines for Java applications deployed to Kubernetes In Detail Imagine creating and testing Java EE applications on Apache Tomcat Server or Wildfly Application server in minutes along with deploying and managing Java applications swiftly. Sounds too good to be true? But you have a reason to cheer as such scenarios are only possible by leveraging Docker and Kubernetes. This book will start by introducing Docker and delve deep into its networking and persistent storage concepts. You will then proceed to learn how to refactor monolith application into separate services by building an application and then packaging it into Docker containers. Next, you will create an image containing Java Enterprise Application and later run it using Docker. Moving on, the book will focus on Kubernetes and its features and you will learn to deploy a Java application to Kubernetes using Maven and monitor a Java application in production. By the end of the book, you will get hands-on with some more advanced topics to further extend your knowledge about Docker and Kubernetes. Style and approach An easy-to-follow, practical guide that will help Java developers develop, deploy, and manage Java applications efficiently.

This book is aimed at system administrators/architects or DBAs who want to learn more about how to grow their current infrastructure to support larger traffic. Before beginning with this book, we expect you to be well-practiced with MySQL/MariaDB for common usage. You will be able to get a grasp quickly if you are comfortable with learning and building large infrastructures for MariaDB using Linux.

Find bottlenecks, identify the proper algorithm to use, optimize performance, and create really efficient Rust applications Key Features Understand common performance pitfalls and improve the performance of your applications. Get to grips with parallel programming and multithreading with Rust. Learn metaprogramming in Rust. Book Description At times, it is difficult to get the best performance out of Rust. This book teaches you to optimize the speed of your Rust code to the level of languages such as C/C++. You'll understand and fix common pitfalls, learn how to improve your productivity by using metaprogramming, and speed up your code by concurrently executing parts of it safely and easily. You will master the features of the language which will make you stand out and use them to really improve the efficiency of your algorithms The book begins with a gentle introduction to help you identify bottlenecks when programming in Rust. We highlight common performance pitfalls, along with strategies to detect and resolve these issues early. We move on to mastering Rust's type system, which will enable us to create impressive optimizations in both performance and safety at compile time. You will then learn how to effectively manage memory in Rust, mastering the borrow checker. We move on to measuring performance and you will see how this affects the way you write code. Moving ahead, you will perform metaprogramming in Rust to boost the performance of your code and your productivity. You will finally learn parallel programming in Rust, which enables efficient and faster execution by using multithreading and asynchronous programming. What you will learn Master tips and tricks to make your code faster. Learn how to identify bottlenecks in your Rust applications Discover how to profile your Rust software. Understand the type system to create compile-time optimizations. Master the borrow checker . Learn metaprogramming in Rust to avoid boilerplate code. Discover multithreading and work stealing in Rust. Understand asynchronous programming in Rust. Who this book is for This book is for Rust developers keen to improve the speed of their code or simply to take their skills to the next level.

This is the latest edition of the book that application developers worldwide have used to master MySQL]€]now updated for MySQL 8 and beyond. As you would expect, this book shows how to code all the essential SQL statements for working with a MySQL database. You'll use these statements every day to have MySQL do more of your work for you. But beyond that, it shows how to work with classic MySQL features that take you to new level, such as summary queries, subqueries, functions, views, transactions, stored procedures, triggers, and security. It shows how to take advantage of newer MySQL features such as window functions, Common Table Expressions (CTE), and roles for database security. It shows how to design a database, including how to use MySQL Workbench to create and implement the design. It even presents a starting set of skills for a database administrator (DBA) if you're interested in that career path or if you need to be your own DBA. In short, it's a must-have guide for anyone who works with MySQL, beginning and experienced developers alike.

Mitigate the risks involved in migrating away from a proprietary database platform toward MariaDB's open source database engine. This book will help you assess the risks and the work involved, and ensure

a successful migration. Migrating to MariaDB describes the process and lessons learned during a migration from a proprietary database management engine to the MariaDB open source solution. The book discusses the drivers for making the decision and change, walking you through all aspects of the process from evaluating the licensing, navigating the pitfalls and hurdles of a migration, through to final implementation on the new platform. The book highlights the cost-effectiveness of MariaDB and how the licensing worries are simplified in comparison to running on a proprietary platform. You'll learn to do your own risk assessment, to identify database and application code that may need to be modified or re-implemented, and to identify MariaDB features to provide the security and failover protection needed by corporate customers. Let the author's experience in migrating a financial firm to MariaDB inform your own efforts, helping you to develop a road map for both technical and political success within your own organization as you migrate away from proprietary lock-in toward MariaDB's open source solution. What You'll Learn Evaluate and compare licensing costs between proprietary databases and MariaDB Perform a proper risk assessment to inform your planning and execution of the migration Build a migration road map from the book's example that is specific to your situation Make needed application changes and migrate data to the MariaDB open source database engine Who This Book Is For Technical professionals (including database administrators, programmers, and technical management) who are interested in migrating away from a proprietary database platform toward MariaDB's open source database engine and need to assess the risks and the work involved

With special focus on the next major release of MySQL, this resource provides a solid framework for anyone new to MySQL or transitioning from another database platform, as well as experience MySQL administrators. The high-profile author duo provides essential coverage of the fundamentals of MySQL database management—including MySQL's unique approach to basic database features and functions—as well as coverage of SQL queries, data and index types, stored procedure and functions, triggers and views, and transactions. They also present comprehensive coverage of such topics as MySQL server tuning, managing storage engines, caching, backup and recovery, managing users, index tuning, database and performance monitoring, security, and more.

Over 90 practical, actionable recipes to automate, test, and manage your infrastructure quickly and effectively About This Book Bring down your delivery timeline from days to hours by treating your server configurations and VMs as code, just like you would with software code. Take your existing knowledge and skill set with your existing tools (Puppet, Chef, or Docker) to the next level and solve IT infrastructure challenges. Use practical recipes to use code to provision and deploy servers and applications and have greater control of your infrastructure. Who This Book Is For This book is for DevOps engineers and developers working in cross-functional teams or operations and would now switch to IAC to manage complex infrastructures. What You Will Learn Provision local and remote development environments with Vagrant Automate production infrastructures with Terraform, Ansible and Cloud-init on AWS, OpenStack, Google Cloud, Digital Ocean, and more Manage and test automated systems using Chef and Puppet Build, ship, and debug optimized Docker containers Explore the best practices to automate and test everything from cloud infrastructures to operating system configuration In Detail Infrastructure as Code (IAC) is a key aspect of the DevOps movement, and this book will show you how to transform the way you work with your infrastructure—by treating it as software. This book is dedicated to helping you discover the essentials of infrastructure automation and its related practices; the over 90 organized practical solutions will demonstrate how to work with some of the very best tools and cloud solutions. You will learn how to deploy repeatable infrastructures and services on AWS, OpenStack, Google Cloud, and Digital Ocean. You will see both Ansible and Terraform in action, manipulate the best bits from cloud-init to easily bootstrap instances, and simulate consistent environments locally or remotely using Vagrant. You will discover how to automate and test a range of system tasks using Chef or Puppet. You will also build, test, and debug various Docker containers having developers' interests in mind. This book will help you to use the right tools, techniques, and approaches to deliver working solutions for today's modern infrastructure challenges. Style and approach This is a recipe-based book that allows you to venture into some of the most cutting-edge practices and techniques about IAC and solve immediate problems when trying to implement them.

If you're a programmer new to databases—or just new to MySQL and its community-driven variant, MariaDB—you've found the perfect introduction. This hands-on guide provides an easy, step-by-step approach to installing, using, and maintaining these popular relational database engines. Author Russell Dyer, Curriculum Manager at MariaDB and former editor of the MySQL Knowledge Base, takes you through database design and the basics of data management and manipulation, using real-world examples and many practical tips. Exercises and review questions help you practice what you've just learned. Create and alter MySQL tables and specify fields and columns within them Learn how to insert, select, update, delete, join, and subquery data, using practical examples Use built-in string functions to find, extract, format, and convert text from columns Learn functions for mathematical or statistical calculations, and for formatting date and time values Perform administrative duties such as managing user accounts, backing up databases, and importing large amounts of data Use APIs to connect and query MySQL and MariaDB with PHP and other languages

A guide to using MySQL covers such topics as accessing and manipulating data, managing security, importing and exporting data, and issuing SQL statements.

This comprehensive reference guide offers useful pointers for advanced use of SQL and describes the bugs and workarounds involved in compiling MySQL for every system.

Unleash the power of XMPP in order to build exciting, real-time, federated applications based on open standards in a secure and highly scalable fashion About This Book Learn about the fundamentals of XMPP and be able to work with the core functionality both server-side and in the browser Build a simple 1-to-1 chat (the "Hello World" of XMPP), explore multi-user chat, publish subscribe systems, and work with a decentralized social network Author Lloyd Watkins is a member of the XMPP standards committee Who This Book Is For If you want to learn about the fundamentals of XMPP, be able to work with the core functionality both server-side and in the browser then this book is for you.No knowledge of XMPP is required, or of TCP/IP networking. It's important that you already know how to build applications of some form, and are looking get a better understanding of how to implement XMPP for one or more of its many uses. You should be interested in the decentralized web, know HTML, and likely know JavaScript and NodeJS. You will probably know JSON, and hopefully XML (this is the native output of XMPP). What You Will Learn Install and configure an XMPP server and use it to connect from a traditional desktop client and send a message Build a simple server-side application that will respond to messages from our logged in desktop client Install and run XMPP-FTW, connect to the server from the browser, and handle incoming/outgoing messages Connect to a multi-user chat room, send/receive stanzas, add a room password, join a protected room, set the room's subject, and change a user's affiliation Get to grips with the publish-subscribe extension of XMPP and use it to build a pusher system that can make any website real-time Build a simple XMPP component and create an extension for XMPP-FTW that allows you to use your own custom

format Build an XMPP version of the classic game “Pong” In Detail XMPP (eXtensible Messaging and Presence Protocol) is a messaging protocol that enables communication between two or more devices via the Internet. With this book, developers will learn about the fundamentals of XMPP, be able to work with the core functionality both server-side and in the browser, as well as starting to explore several of the protocol extensions. You will not only have a solid grasp of XMPP and how it works, but will also be able to use the protocol to build real-world applications that utilize the power of XMPP. By the end of this book, you will know more about networking applications in general, and have a good understanding of how to extend XMPP, as well as using it in sample applications. Style and approach Through a number of hands-on projects, this book shows you how to build usable applications that highlights a feature of XMPP.

Quickly get up to speed with MariaDB—the leading, drop-in replacement for MySQL, through this practical tutorial About This Book Get to know the basic SQL queries so you can quickly start using MariaDB Take control of your data through the advanced features of MariaDB Exploit the full potential of MariaDB's exclusive features through quick, practical examples Who This Book Is For If you don't know the SQL language, but you want to quickly jump into the SQL world and learn how to use MariaDB, or if you already know how to use MySQL but you want to go further, then this book is ideal for you. What You Will Learn Install and configure MariaDB Create databases, tables, and indexes Import and export data from and to external files Work with views and virtual columns Create, read, update, and delete records in your database Use dynamic columns Set up a powerful full-text search system Access your external data from MariaDB through the CONNECT engine In Detail This book will take you through all the nitty-gritty parts of MariaDB, right from the creation of your database all the way to using MariaDB's advanced features. At the very beginning, we show you the basics, that is, how to install MariaDB. Then, we walk you through the databases and tables of MariaDB, and introduce SQL in MariaDB. You will learn about all the features that have been added in MariaDB but are absent in MySQL. Moving on, you'll learn to import and export data, views, virtual columns, and dynamic columns in MariaDB. Then, you'll get to grips with full-text searches and queries in MariaDb. You'll also be familiarized with the CONNECT storage engine. At the end of the book, you'll be introduced to the community of MariaDB. Style and approach This is a complete guide that uses concrete examples to help you understand and exploit the full potential of MariaDB.

Server bottlenecks and failures are a fact of life in any database deployment, but they don't have to bring everything to a halt. This practical book explains replication, cluster, and monitoring features that can help protect your MySQL system from outages, whether it's running on hardware, virtual machines, or in the cloud. Written by engineers who designed many of the tools covered, this book reveals undocumented or hard-to-find aspects of MySQL reliability and high availability—knowledge that's essential for any organization using this database system. This second edition describes extensive changes to MySQL tools. Versions up to 5.5 are covered, along with several 5.6 features. Learn replication fundamentals, including use of the binary log and MySQL Replicant Library Handle failing components through redundancy Scale out to manage read-load increases, and use data sharding to handle large databases and write-load increases Store and replicate data on individual nodes with MySQL Cluster Monitor database activity and performance, and major operating system parameters Keep track of masters and slaves, and deal with failures and restarts, corruption, and other incidents Examine tools including MySQL Enterprise Monitor, MySQL Utilities, and GTIDs

This two-volume book presents an unusually diverse selection of research papers, covering all major topics in the fields of information and communication technologies and related sciences. It provides a wide-angle snapshot of current themes in information and power engineering, pursuing a cross-disciplinary approach to do so. The book gathers revised contributions that were presented at the 2018 International Conference: Sciences of Electronics, Technologies of Information and Telecommunication (SETIT'18), held on 20–22 December 2018 in Hammamet, Tunisia. This eighth installment of the event attracted a wealth of submissions, and the papers presented here were selected by a committee of experts and underwent additional, painstaking revision. Topics covered include: · Information Processing · Human-Machine Interaction · Computer Science · Telecommunications and Networks · Signal Processing · Electronics · Image and Video This broad-scoped approach is becoming increasingly popular in scientific publishing. Its aim is to encourage scholars and professionals to overcome disciplinary barriers, as demanded by current trends in the industry and in the consumer market, which are rapidly leading toward a convergence of data-driven applications, computation, telecommunication, and energy awareness. Given its coverage, the book will benefit graduate students, researchers and practitioners who need to keep up with the latest technological advances.

Create high availability clusters to enhance system performance using CentOS 7 About This Book Master the concepts of high performance and high availability to eliminate performance bottlenecks Maximize the uptime of services running in a CentOS 7 cluster A step-by-step guide that will provide knowledge of methods and approaches to optimize the performance of CentOS clusters Who This Book Is For This book is targeted at system administrators: those who want a detailed, step-by-step guide to learn how to set up a high-availability CentOS 7 cluster, and those who are looking for a reference book to help them learn or refresh the necessary skills to ensure their systems and respective resources are utilized optimally. No previous knowledge of high-availability systems is needed, though the reader is expected to have at least some degree of familiarity with any spin-off of the Fedora family of Linux distributions, preferably CentOS. What You Will Learn Install a CentOS 7 cluster and network infrastructure Configure firewall, networking, and clustering services and settings Set up and test a HAC (high-availability cluster) to host an Apache web server and a MariaDB database server Monitor performance and availability Identify bottlenecks and troubleshoot issues Improve performance and ensure high availability In Detail CentOS is the enterprise level Linux OS, which is 100% binary compatible to Red Hat Enterprise Linux (RHEL). It acts as a free alternative to RedHat's commercial Linux offering, with only a change in the branding. A high performance cluster consists in a group of computers that work together as one set parallel, hence minimizing or eliminating the downtime of critical services and enhancing the performance of the application. Starting with the basic principles of clustering, you will learn the necessary steps to install a cluster with two CentOS 7 servers. We will then set up and configure the basic required network infrastructure and clustering services. Further, you will learn how to take a proactive approach to the split-brain issue by configuring the failover and fencing of the cluster as a whole and the quorum of each node individually. Further, we will be setting up HAC and HPC clusters as a web server and a database server. You will also master the art of monitoring performance and availability, identifying bottlenecks, and exploring troubleshooting techniques. At the end of the book, you'll review performance-tuning techniques for the recently installed cluster, test performance using a payload simulation, and learn the necessary skills to ensure that the systems, and the corresponding resources and services, are being utilized to their best capacity. Style and approach An easy-to-follow and step-by-step guide with hands-on instructions to set up real-world simple cluster scenarios that will start you on the path to building more complex applications on your own.

"With an easy, step-by-step approach, this guide shows beginners how to install, use, and maintain the world's most popular open source database: MySQL. You'll learn through real-world examples and many practical tips, including information on how to improve database performance. Database systems such as MySQL help data handling for organizations large and small handle data, providing robust and efficient access in ways not offered by spreadsheets and other types of data stores. This book is also useful for web developers and programmers interested in adding MySQL to their skill sets. Topics

include: Installation and basic administration ; Introduction to databases and SQL ; Functions, subqueries, and other query enhancements ; Improving database performance ; Accessing MySQL from popular languages" --

How can you help your Drupal website continue to perform at the highest level as it grows to meet demand? This comprehensive guide provides best practices, examples, and in-depth explanations for solving several performance and scalability issues. You'll learn how to apply coding and infrastructure techniques to Drupal internals, application performance, databases, web servers, and performance analysis. Covering Drupal versions 7 and 8, this book is the ideal reference for everything from site deployment to implementing specific technologies such as Varnish, memcache, or Solr. If you have a basic understanding of Drupal and the Linux-Apache-MySQL-PHP (LAMP) stack, you're ready to get started. Establish a performance baseline and define goals for improvement Optimize your website's code and front-end performance Get best and worst practices for customizing Drupal core functionality Apply infrastructure design techniques to launch or expand a site Use tools to configure, monitor, and optimize MySQL performance Employ alternative storage and backend search options as your site grows Tune your web servers through httpd and PHP configuration Monitor services and perform load tests to catch problems before they become critical

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 as it does, with illustrative stories and case studies that demonstrate MySQL's principles in action. With this book, you'll learn how to think in MySQL. Learn the effects of new features in MySQL 5.5, including stored procedures, partitioned databases, triggers, and views Implement improvements in replication, high availability, and clustering Achieve high performance when running MySQL in the cloud Optimize advanced querying features, such as full-text searches Take advantage of modern multi-core CPUs and solid-state disks Explore backup and recovery strategies—including new tools for hot online backups

If you are a system administrator or Linux professional who wants to learn to set up, install, and manage OpenVZ containers on a server to implement OS-level virtualization, then this book is for you. Along with elementary knowledge of Linux programming, you need to have a conceptual understanding of system components and functions.

A practical cookbook, filled with advanced recipes, and plenty of code and commands used for illustration, which will make your learning curve easy and quick. This book is for anyone who wants to learn more about databases in general or MariaDB in particular. Some familiarity with SQL databases is assumed, but the recipes are approachable to almost anyone with basic database skills.

How can you realize MySQL's full power? With High Performance MySQL, you'll learn advanced techniques for everything from choosing the right abstraction layer for databases to designing schemas, indexes, and queries to tuning your server, operating system, and hardware to achieve their 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, features, and tools, this revised edition helps you design a data platform that will scale with your business. You'll learn the latest in cloud-hosted MySQL offerings, best practices for database security, and 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

MariaDB is a database server that offers drop-in replacement functionality for MySQL. Built by some of the original authors of MySQL, with assistance from the broader community of free and open source software developers, MariaDB offers a rich set of feature enhancements to MySQL, including alternate storage engines, server optimizations, and patches. MariaDB Crash Course teaches you all you need to know to be immediately productive with MariaDB. Master trainer Ben Forta introduces all the essentials through a series of quick, easy-to-follow, hands-on lessons. Instead of belaboring database theory and relational design, Forta focuses on teaching solutions for the majority of users who simply want to interact with data. Learn how to: Retrieve and sort data Filter data using comparisons, regular expressions, and full text search Join relational data Create and alter tables Insert, update, and delete data Leverage the power of stored procedures and triggers Use views and cursors Manage transactional processing Create user accounts and manage security via access control

MariaDB High PerformancePackt Publishing Ltd

In this book, you will learn how to build from scratch a criminal records management database system using MariaDB Connector/J. As you know, MariaDB server is a community developed fork of MySQL server. Started by core members of the original MySQL team, MariaDB actively works with outside developers to deliver the most featureful, stable, and sanely licensed open SQL server in the industry. In the first chapter, you will be taught how to create Crime database and its tables. In the second chapter, you will create Suspect table. You will be taught to create Java GUI to view, edit, insert, and delete Suspect table data. This table has eleven columns: suspect_id (primary key), suspect_name, birth_date, case_date, report_date, suspect_status, arrest_date, mother_name, address, telephone, and photo. In the third chapter, you will be taught to create Java GUI to view, edit, insert, and delete Feature_Extraction table data. This table has eight columns: feature_id (primary key), suspect_id (foreign key), feature1, feature2, feature3, feature4, feature5, and feature6. All six fields (except keys) will have a BLOB data type, so that the image of the feature will be directly saved into this table. In the fourth chapter, you will add two tables: Police_Station and Investigator. These two tables will later be joined to Suspect table through another table, File_Case, which will be built in the seventh chapter. The Police_Station has six columns: police_station_id (primary key), location, city, province, telephone, and photo. The Investigator has eight columns: investigator_id (primary key), investigator_name, rank, birth_date, gender, address, telephone, and photo. Here, you will design a Java GUI to display, edit, fill, and delete data in both tables. In the fifth chapter, you will add two tables: Victim and File_Case. The File_Case table will connect four other tables: Suspect, Police_Station, Investigator and Victim. The Victim table has nine columns: victim_id (primary key), victim_name, crime_type, birth_date, crime_date, gender, address, telephone, and photo. The File_Case has seven columns: file_case_id (primary key), suspect_id (foreign key), police_station_id (foreign key), investigator_id (foreign key), victim_id (foreign key), status, and description. Here, you will also design a Java GUI to display, edit, fill, and delete data in both tables.

Take your knowledge of Redis to the next level to build enthralling applications with ease About This Book Detailed explanation on Data structure server with powerful strings, lists, sets, sorted-sets, and hashes Learn to Scale your data with Redis Cluster's distributed setup This is a fast paced practical guide full of screenshots and real work examples to help you get to grips with Redis in no time. Who This Book Is For If you are a software developer with some experience with Redis and would now like to elevate your Redis knowledge and skills even further, then this book is for you. What You Will Learn Choose the right Redis data structure for your problem Understand Redis event-loop and implement your own custom C commands Solve complex workflows with Redis server-side scripting with Lua Configure your Redis instance for optimal memory management Scale your data in a distributed manner with Redis Cluster Improve the stability of your Redis solution using Redis Sentinel Complement your

existing database and NoSQL environment with Redis. Exploit a wide range of features provided by Redis to become a DevOps expert. In Detail Redis is the most popular, open-source, key value data structure server that provides a wide range of capabilities on which multiple platforms can be built. Its fast and flexible data structures give your existing applications an edge in the development environment. This book is a practical guide which aims to help you deep dive into the world of Redis data structure to exploit its excellent features. We start our journey by understanding the need of Redis in brief, followed by an explanation of Advanced key management. Next, you will learn about design patterns, best practices for using Redis in DevOps environment and Docker containerization paradigm in detail. After this, you will understand the concept of scaling with Redis cluster and Redis Sentinel, followed by a thorough explanation of incorporating Redis with NoSQL technologies such as Elasticsearch and MongoDB. At the end of this section, you will be able to develop competent applications using these technologies. You will then explore the message queuing and task management features of Redis and will be able to implement them in your applications. Finally, you will learn how Redis can be used to build real-time data analytic dashboards, for different disparate data streams. Style and approach This is a hands on guide full of easy-to-follow examples, that illustrate important concepts and techniques to solve complex problems with Redis.

This book is a collection of developer code recipes and best practices for persisting data using Spring, particularly Spring Boot. The book is structured around practical recipes, where each recipe discusses a performance case or performance-related case, and almost every recipe has one or more applications. Mainly, when we try to accomplish something (e.g., read some data from the database), there are several approaches to do it, and, in order to choose the best way, you have to know the implied trades-off from a performance perspective. You'll see that in the end, all these penalties slow down the application. Besides presenting the arguments that favor a certain choice, the application is written in Spring Boot style which is quite different than plain Hibernate. Persistence is an important set of techniques and technologies for accessing and using data, and this book demonstrates that data is mobile regardless of specific applications and contexts. In Java development, persistence is a key factor in enterprise, ecommerce, cloud and other transaction-oriented applications. After reading and using this book, you'll have the fundamentals to apply these persistence solutions into your own mission-critical enterprise Java applications that you build using Spring. What You Will Learn Shape *-to-many associations for best performances Effectively exploit Spring Projections (DTO) Learn best practices for batching inserts, updates and deletes Effectively fetch parent and association in a single SELECT Learn how to inspect Persistent Context content Dissect pagination techniques (offset and keyset) Handle queries, locking, schemas, Hibernate types, and more Who This Book Is For Any Spring and Spring Boot developer that wants to squeeze the persistence layer performances.

A guide for MySQL administrators covers such topics as benchmarking, server performance, indexing, queries, hardware optimization, replication, scaling, cloud hosting, and backup and recovery.

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.

Quickly get up to speed with MariaDB—the leading, drop-in replacement for MySQL, through this practical tutorial About This Book • Get to know the basic SQL queries so you can quickly start using MariaDB • Take control of your data through the advanced features of MariaDB • Exploit the full potential of MariaDB's exclusive features through quick, practical examples Who This Book Is For If you don't know the SQL language, but you want to quickly jump into the SQL world and learn how to use MariaDB, or if you already know how to use MySQL but you want to go further, then this book is ideal for you. What You Will Learn • Install and configure MariaDB • Create databases, tables, and indexes • Import and export data from and to external files • Work with views and virtual columns • Create, read, update, and delete records in your database • Use dynamic columns • Set up a powerful full-text search system • Access your external data from MariaDB through the CONNECT engine In Detail This book will take you through all the nitty-gritty parts of MariaDB, right from the creation of your database all the way to using MariaDB's advanced features. At the very beginning, we show you the basics, that is, how to install MariaDB. Then, we walk you through the databases and tables of MariaDB, and introduce SQL in MariaDB. You will learn about all the features that have been added in MariaDB but are absent in MySQL. Moving on, you'll learn to import and export data, views, virtual columns, and dynamic columns in MariaDB. Then, you'll get to grips with full-text searches and queries in MariaDB. You'll also be familiarized with the CONNECT storage engine. At the end of the book, you'll be introduced to the community of MariaDB. Style and approach This is a complete guide that uses concrete examples to help you understand and exploit the full potential of MariaDB.

High Performance MySQL is the definitive guide to building fast, reliable systems with MySQL. Written by noted experts with years of real-world experience building very large systems, this book covers every aspect of MySQL performance in detail, and focuses on robustness, security, and data integrity. High Performance MySQL teaches you advanced techniques in depth so you can bring out MySQL's full power. Learn how to design schemas, indexes, queries and advanced MySQL features for maximum performance, and get detailed guidance for tuning your MySQL server, operating system, and hardware to their fullest potential. You'll also learn practical, safe, high-performance ways to scale your applications with replication, load balancing, high availability, and failover. This second edition is completely revised and greatly expanded, with deeper coverage in all areas. Major additions include: Emphasis throughout on both performance and reliability Thorough coverage of storage engines, including in-depth tuning and optimizations for the InnoDB storage engine Effects of new features in MySQL 5.0 and 5.1, including stored procedures, partitioned databases, triggers, and views A detailed discussion on how to build very large, highly scalable systems with MySQL New options for backups and replication Optimization of advanced querying features, such as full-text searches Four new appendices The book also includes chapters on benchmarking, profiling, backups, security, and tools and techniques to help you measure, monitor, and manage your MySQL installations.

This book is intended for intermediate users who want to learn how to administrate a MariaDB server or a set of servers. It is aimed at MariaDB users, and hence working knowledge of MariaDB is a prerequisite.

Businesses, right from small-sized to large giants are using MySQL to access and manipulate their enterprise database across the world. Advanced MySQL 8 will be your one-stop destination in how to implement and manage large-scale MySQL distributed clusters to run high-volume websites, enterprise systems, and packaged software.

As users come to depend on MySQL, they find that they have to deal with issues of reliability, scalability, and performance--issues that are not well documented but are critical to a smoothly functioning site. This book is an insider's guide to these little understood topics. Author Jeremy Zawodny has managed large numbers of MySQL servers for mission-critical work at Yahoo!, maintained years of contacts with the MySQL AB team, and presents regularly at conferences. Jeremy and Derek have spent months experimenting, interviewing major users of MySQL, talking to MySQL AB, benchmarking, and writing some of their own tools in order to produce the information in this book. In High Performance MySQL you will learn about MySQL indexing and optimization in depth so you can make better use of these key features. You will learn practical replication, backup, and load-balancing strategies with information that goes beyond available tools to discuss their effects in real-life environments. And you'll learn the supporting techniques you need to carry out these tasks, including advanced configuration, benchmarking, and investigating logs. Topics include: A review of configuration and setup options Storage engines and table types Benchmarking Indexes Query Optimization Application Design Server Performance Replication Load-balancing Backup and Recovery Security

This book constitutes the refereed proceedings of the 11th International Conference entitled Beyond Databases, Architectures and Structures, BDAS 2015, held in Ustro?, Poland, in May 2015. This book consists of 53 carefully revised selected papers that are assigned to 8 thematic groups: database architectures and performance; data integration, storage and data warehousing; ontologies and semantic web; artificial intelligence, data mining and knowledge discovery; image analysis and multimedia mining; spatial data analysis; database systems development; application of database systems.

If you are a software developer or administrator who wishes to create simple, reusable environments using Vagrant, this book is the perfect choice for you. Whether you are a system administrator with extensive experience in virtualization or a developer wishing to create development scripts for cloud deployment, you will find easy-to-follow recipes and techniques in this book that will allow you to create robust and reproducible virtual environments.

A practical guide to administer, monitor and replicate your PostgreSQL 11 database Key Features Study and apply the newly introduced features in PostgreSQL 11 Tackle any problem in PostgreSQL 11 administration and management Catch up on expert techniques for monitoring, fine-tuning, and securing your database Book Description PostgreSQL is a powerful, open source database management system with an enviable reputation for high performance and stability. With many new features in its arsenal, PostgreSQL 11 allows you to scale up your PostgreSQL infrastructure. This book takes a step-by-step, recipe-based approach to effective PostgreSQL administration. The book will introduce you to new features such as logical replication, native table partitioning, additional query parallelism, and much more to help you to understand and control, crash recovery and plan backups. You will learn how to tackle a variety of problems and pain points for any database administrator such as creating tables, managing views, improving performance, and securing your database. As you make steady progress, the book will draw attention to important topics such as monitoring roles, backup, and recovery of your PostgreSQL 11 database to help you understand roles and produce a summary of log files, ensuring high availability, concurrency, and replication. By the end of this book, you will have the necessary knowledge to manage your PostgreSQL 11 database efficiently. What you will learn Troubleshoot open source PostgreSQL version 11 on various platforms Deploy best practices for planning and designing live databases Select and implement robust backup and recovery techniques in PostgreSQL 11 Use pgAdmin or OmniDB to perform database administrator (DBA) tasks Adopt efficient replication and high availability techniques in PostgreSQL Improve the performance of your PostgreSQL solution Who this book is for This book is designed for database administrators, data architects, database developers, or anyone with an interest in planning and running live production databases using PostgreSQL 11. It is also ideal if you're looking for hands-on solutions to any problem associated with PostgreSQL 11 administration. Some experience with handling PostgreSQL databases will be beneficial

This book explains relational theory in practice, and demonstrates through two projects how you can apply it to your use of MariaDB and SQL Server databases. This book covers the important requirements of teaching databases with a practical and progressive perspective. This book offers the straightforward, practical answers you need to help you do your job. This hands-on tutorial/reference/guide to MariaDB and SQL Server is not only perfect for students and beginners, but it also works for experienced developers who aren't getting the most from MariaDB and SQL Server. As you would expect, this book shows how to build from scratch two different databases: MariaDB and SQL Server using Java. In designing a GUI and as an IDE, you will make use of the NetBeans tool. In chapter one, you will learn the basics of cryptography using Java. Here, you will learn how to write a Java program to count Hash, MAC (Message Authentication Code), store keys in a KeyStore, generate PrivateKey and PublicKey, encrypt / decrypt data, and generate and verify digital prints. You will also learn how to create and store salt passwords and verify them. In chapter two, you will create a PostgreSQL database, named Bank, and its tables. In chapter three, you will create a Login table. In this case, you will see how to create a Java GUI using NetBeans to implement it. In addition to the Login table, in this chapter you will also create a Client table. In the case of the Client table, you will learn how to generate and save public and private keys into a database. You will also learn how to encrypt / decrypt data and save the results into a database. In chapter four, you will create an Account table. This account table has the following ten fields: account_id (primary key), client_id (primarykey), account_number, account_date, account_type, plain_balance, cipher_balance, decipher_balance, digital_signature, and signature_verification. In this case, you will learn how to implement generating and verifying digital prints and storing the results into a database. In chapter five, you create a table named Client_Data, which has seven columns: client_data_id (primary key), account_id (primary_key), birth_date, address, mother_name, telephone, and photo_path. In chapter six, you will be taught how to create a SQL Server database, named Crime, and its tables. In chapter seven, you will be taught how to extract image features, utilizing BufferedImage class, in Java GUI. In chapter eight, you will be taught to create Java GUI to view, edit, insert, and delete Suspect table data. This table has eleven columns: suspect_id (primary key), suspect_name, birth_date, case_date, report_date, suspect_status, arrest_date, mother_name, address, telephone, and photo. In chapter nine, you will be taught to create Java GUI to view, edit, insert, and delete Feature_Extraction table data. This table has eight columns: feature_id (primary key), suspect_id (foreign key), feature1, feature2, feature3, feature4, feature5, and feature6. In chapter ten, you will add two tables: Police_Station and Investigator. These two tables will later be joined to Suspect table through another table, File_Case, which will be built in the seventh chapter. The Police_Station has six columns: police_station_id (primary key), location, city, province, telephone, and photo. The Investigator has eight columns: investigator_id (primary key), investigator_name, rank, birth_date, gender, address, telephone, and photo. Here, you will design a Java GUI to display, edit, fill, and delete data in both tables. In chapter eleven, you will add two tables: Victim and File_Case. The File_Case table will connect four other tables: Suspect, Police_Station, Investigator and Victim. The Victim table has nine columns: victim_id (primary key), victim_name, crime_type, birth_date, crime_date, gender, address, telephone, and photo. The File_Case has seven columns: file_case_id (primary key), suspect_id (foreign key), police_station_id (foreign key), investigator_id (foreign key), victim_id (foreign key), status, and description. Here, you will also design a Java GUI to display, edit, fill, and delete data in both tables. Finally, this book is hopefully useful and can improve database programming skills for every Java/MariaDB/SQL Server programmer.

[Copyright: 9027b311e94e23b43fee36cc57997d83](https://www.pdfdrive.com/mariadb-high-performance-ebook.html)