

Mariadb High Performance

Leverage the power of PostgreSQL 10 to design, administer and maintain a high-performance database solution

Key Features Obtain optimal PostgreSQL 10 database performance, ranging from initial design to routine maintenance

Fine tune the performance of your queries and avoid the common pitfalls that can slow your system down

Contains tips and tricks on scaling successful database installations, and ensuring a highly available PostgreSQL solution

Book Description PostgreSQL database servers have a common set of problems that they encounter as their usage gets heavier and requirements get more demanding. Peek into the future of your PostgreSQL 10 database's problems today. Know the warning signs to look for and how to avoid the most common issues before they even happen. Surprisingly, most PostgreSQL database applications evolve in the same way--choose the right hardware, tune the operating system and server memory use, optimize queries against the database and CPUs with the right indexes, and monitor every layer, from hardware to queries, using tools from inside and outside PostgreSQL. Also, using monitoring insight, PostgreSQL database applications continuously rework the design and configuration. On reaching the limits of a single server, they break things up; connection pooling, caching, partitioning, replication, and parallel queries can all help handle increasing database workloads. By the end of this book, you will have all the knowledge you need to design, run, and manage your PostgreSQL solution while ensuring high performance and high availability

What you will learn Learn best practices for scaling PostgreSQL 10 installations

Discover the best hardware for developing high-performance PostgreSQL applications

Benchmark your whole system - from hardware to application

Learn by real examples how server parameters impact performance

Discover PostgreSQL 10 features for partitioning and parallel query

Monitor your server, both inside and outside the database

Design and implement a good replication system on PostgreSQL 10

Who this book is for This book is designed for database administrators and PostgreSQL architects who already use or plan to exploit the features of PostgreSQL 10 to design and maintain a high-performance PostgreSQL database. A working knowledge of SQL, and some experience with PostgreSQL will be helpful in getting the most out of this book.

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.

Build interactive, database-driven websites with PHP 7, MySQL 8, and MariaDB. The focus of this book is on getting you up and running as quickly as possible with real-world applications. In the first two chapters, you will set up your development and testing environment, and then build your first PHP and

MariaDB or MySQL database-driven website. You will then increase its sophistication, security, and functionality throughout the course of the book. The PHP required is taught in context within each project so you can quickly learn how PHP integrates with MariaDB and MySQL to create powerful database-driven websites. Each project is fully illustrated, so you will see clearly what you are building as you create your own database-driven website. You will build a form for registering users, and then build an interface so that an administrator can view and administer the user database. You will create a message board for users and a method for emailing them. You will also learn the best practices for ensuring that your website databases are secure. Later chapters describe how to create a product catalog, and a simple e-commerce site. You will also discover how to migrate a database to a remote host. The final chapter will demonstrate the advantages of migrating to Oracle's MySQL 8. You will be shown step by step migration directions along with a demonstration of the tools available in SQL Workbench. Because you are building the interactive pages yourself, you will know exactly how MySQL, MariaDB, and PHP all work together, and you will be able to add database interactivity to your own websites with ease. What You Will Learn Build a secure database-driven website using PHP 7, MySQL 8, and MariaDB Create a product catalog Write a message board Move towards e-commerce Employ security and validation measures Migrate to Oracle's MySQL 8 Server platform Who This Book Is For Web developers with HTML, CSS and a limited Bootstrap experience. Readers need little to no prior experience with PHP and MySQL.

Beginning and experienced programmers will use this comprehensive guide to persistent memory programming. You will understand how persistent memory brings together several new software/hardware requirements, and offers great promise for better performance and faster application startup times—a huge leap forward in byte-addressable capacity compared with current DRAM offerings. This revolutionary new technology gives applications significant performance and capacity improvements over existing technologies. It requires a new way of thinking and developing, which makes this highly disruptive to the IT/computing industry. The full spectrum of industry sectors that will benefit from this technology include, but are not limited to, in-memory and traditional databases, AI, analytics, HPC, virtualization, and big data. Programming Persistent Memory describes the technology and why it is exciting the industry. It covers the operating system and hardware requirements as well as how to create development environments using emulated or real persistent memory hardware. The book explains fundamental concepts; provides an introduction to persistent memory programming APIs for C, C++, JavaScript, and other languages; discusses RMDA with persistent memory; reviews security features; and presents many examples. Source code and examples that you can run on your own systems are included. What You'll Learn Understand what persistent memory is, what it does, and the value it brings to the industry Become familiar with the operating system and hardware requirements to use persistent memory Know the fundamentals of persistent memory programming: why it is different from current programming methods, and

what developers need to keep in mind when programming for persistence Look at persistent memory application development by example using the Persistent Memory Development Kit (PMDK) Design and optimize data structures for persistent memory Study how real-world applications are modified to leverage persistent memory Utilize the tools available for persistent memory programming, application performance profiling, and debugging Who This Book Is For C, C++, Java, and Python developers, but will also be useful to software, cloud, and hardware architects across a broad spectrum of sectors, including cloud service providers, independent software vendors, high performance compute, artificial intelligence, data analytics, big data, etc.

MariaDB Cookbook

Vagrant Virtual Development Environment Cookbook

High Performance MySQL

Getting Started with MariaDB

Proceedings of the 8th International Conference on Sciences of Electronics, Technologies of Information and Telecommunications (SETIT'18), Vol.1

What To Do When Queries Don't Work

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 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

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.

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.

Fast and Scalable Designs

Mastering Redis

Enterprise-level Features for Scalability and Availability

PostgreSQL 10 High Performance

MySQL 8 Administrator's Guide

Infrastructure as Code (IAC) Cookbook

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.

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

MariaDB High PerformancePackt Publishing Ltd

Run queries and analysis on big data clusters across relational and non relational databases KEY FEATURES ? Connect to Hadoop, Azure, Spark, Oracle, Teradata, Cassandra, MongoDB, CosmosDB, MySQL, PostgreSQL, MariaDB, and SAP HANA. ? Numerous techniques on

how to query data and troubleshoot Polybase for better data analytics. ? Exclusive coverage on Azure Synapse Analytics and building Big Data clusters. DESCRIPTION This book brings exciting coverage on establishing and managing data virtualization using polybase. This book teaches how to configure polybase on almost all relational and nonrelational databases. You will learn to set up the test environment for any tool or software instantly without hassle. You will practice how to design and build some of the high performing data warehousing solutions and that too in a few minutes of time. You will almost become an expert in connecting to all databases including hadoop, cassandra, MySQL, PostgreSQL, MariaDB and Oracle database. This book also brings exclusive coverage on how to build data clusters on Azure and using Azure Synapse Analytics. By the end of this book, you just don't administer the polybase for managing big data clusters but rather you learn to optimize and boost the performance for enabling data analytics and ease of data accessibility. WHAT YOU WILL LEARN ? Learn to configure Polybase and process Transact SQL queries with ease. ? Create a Docker container with SQL Server 2019 on Windows and Polybase. ? Establish SQL Server instance with any other software or tool using Polybase ? Connect with Cassandra, MongoDB, MySQL, PostgreSQL, MariaDB, and IBM DB2. WHO THIS BOOK IS FOR This book is for database developers and administrators familiar with the SQL language and command prompt. Managers and decision-makers will also find this book useful. No prior knowledge of any other technology or language is required. TABLE OF CONTENTS 1. What is Data Virtualization (Polybase) 2. History of Polybase 3. Polybase current state 4. Differences with other technologies 5. Usage 6. Future 7. SQL Server 8. Hadoop Cloudera and Hortonworks 9. Windows Azure Storage Blob 10. Spark 11. From Azure Synapse Analytics 12. From Big Data Clusters 13. Oracle 14. Teradata 15. Cassandra 16. MongoDB 17. CosmosDB 18. MySQL 19. PostgreSQL 20. MariaDB 21. SAP HANA 22. IBM DB2 23. Excel

MySQL Reference Manual

Optimization, Backups, Replication, and More

MariaDB High Performance

Tools for Building Robust Data Centers

IBM Spectrum LSF Suite: Installation Best Practices Guide

Hibernate Tips

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.

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 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.

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.

Programming Persistent Memory

Conquering Big Data with High Performance Computing

Optimization, Backups, and Replication

Discover the full potential of MySQL and ensure high performance of your database

Learning MySQL and MariaDB

Effective guide to administering high-performance MySQL 8 solutions

"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" --

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 through 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 IBM® Redpaper publication describes IBM Spectrum® LSF® Suite best practices installation topics, application checks for workload management, and high availability configurations by using theoretical knowledge and hands-on exercises. These findings are documented by way of sample scenarios. This publication addresses topics for sellers, IT architects, IT specialists, and anyone who wants to implement and manage a high-performing workload management solution with LSF. Moreover, this guide provides

documentation to transfer how-to-skills to the technical teams, and solution guidance to the sales team. This publication compliments documentation that is available at IBM Knowledge Center, and aligns with educational materials that are provided by IBM Systems. 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.

CentOS High Performance

Object Oriented Modular Programming using HTML5, CSS3, JavaScript, XML, JSON, and MySQL

Efficient MySQL Performance

MariaDB Essentials

MySQL High Availability

When you use Hibernate in your projects, you quickly recognize that you need to do more than just add @Entity annotations to your domain model classes. Real-world applications often require advanced mappings, complex queries, custom data types and caching. Hibernate can do all of that. You just have to know which annotations and APIs you need to use. *Hibernate Tips - More than 70 solutions to common Hibernate problems* shows you how to

efficiently implement your persistence layer with Hibernate's basic and advanced features. Each Hibernate Tip consists of one or more code samples and an easy to follow step-by-step explanation. You can also download an example project with executable test cases for each Hibernate Tip. Throughout this book, you will get more than 70 ready-to-use solutions that show you how to:

- Define standard mappings for basic attributes and entity associations.
- Implement your own attribute mappings and support custom data types.
- Use Hibernate's Java 8 support and other proprietary features.
- Read data from the database with JPQL, Criteria API, and native SQL queries.
- Call stored procedures and database functions.

This book is for developers who are already working with Hibernate and who are looking for solutions for their current development tasks. It's not a book for beginners who are looking for extensive descriptions of Hibernate's general concepts. The tips are designed as self-contained recipes which provide a specific solution and can be accessed when needed. Most of them contain links to related tips which you can follow if you want to dive deeper into a topic or need a slightly different solution. There is no need to read the tips in a specific order. Feel free to read the book from cover to cover or to just pick the tips that help you in your current project.

This IBM® Redbooks® publication demonstrates and documents that IBM Power Systems™ high-performance computing and technical computing solutions deliver faster time to value with powerful solutions. Configurable into highly scalable Linux clusters, Power Systems offer extreme performance for demanding workloads such as genomics, finance, computational chemistry, oil and gas exploration, and high-performance data analytics. This book delivers a high-performance computing solution implemented on the IBM Power System S822LC. The solution delivers high application performance and throughput based on its built-for-big-data architecture that incorporates IBM POWER8® processors, tightly coupled Field Programmable Gate Arrays (FPGAs) and accelerators, and faster I/O by using Coherent Accelerator Processor Interface (CAPI). This solution is ideal for clients that need more processing power while simultaneously increasing workload density and reducing datacenter floor space requirements. The Power S822LC offers a modular design to scale from a single rack to hundreds, simplicity of ordering, and a strong innovation roadmap for graphics processing units (GPUs). This publication is targeted toward technical professionals (consultants, technical support staff, IT Architects, and IT Specialists) responsible for delivering cost effective high-performance computing (HPC) solutions that help uncover insights from their data so they can optimize business results, product development, and scientific discoveries

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. 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.

MySQL Troubleshooting

Practical XMPP

Discovering and Improving a Great Database

Documentation from the Source

Implementing an IBM High-Performance Computing Solution on IBM Power System S822LC

Administer Big Data, SQL Queries and Data Accessibility Across Hadoop, Azure, Spark, Cassandra, MongoDB, CosmosDB, MySQL and PostgreSQL (English Edition)

A practical, hands-on, beginner-friendly guide to installing and using MariaDB. Getting Started with MariaDB is for anyone who wants to learn more about databases in general or MariaDB in particular. No prior database experience is required. It is assumed that you have basic knowledge of software installation, editing files with a text editor, and using the command line and terminal.

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.

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.

You'll find several books on basic or advanced MySQL performance, but nothing in between. That's because explaining MySQL performance without addressing its complexity is difficult. This practical book bridges the gap by teaching software engineers mid-level MySQL knowledge beyond the fundamentals, but well shy of deep-level internals required by database administrators (DBAs).

Daniel Nichter shows you how to apply the best practices and techniques that directly affect MySQL performance. You'll learn how to improve performance by analyzing query execution, indexing for common SQL clauses and table joins, optimizing data access, and understanding the most important MySQL metrics. You'll also discover how replication, transactions, row locking, and the cloud influence MySQL performance. Understand why query response time is the North Star of MySQL performance Learn query metrics in detail, including aggregation, reporting, and analysis See how to index effectively for common SQL clauses and table joins Explore the most important server metrics and what they reveal about performance Dive into transactions and row locking to gain deep, actionable insight Achieve remarkable MySQL performance at any scale

Docker and Kubernetes for Java Developers

A Comprehensive Guide for Developers

Practical PHP 7, MySQL 8, and MariaDB Website Databases

High Performance Drupal

Advanced MySQL 8

A Simplified Approach to Developing Database-Driven Websites

Learn about the family of products that make up the MariaDB ecosystem so that you can make the correct choices in applying MariaDB to your own business problems. This video introduces the full suite of scalability and availability features that combine to make MariaDB a compelling option to large enterprises such as major financial institutions and some of the world's largest ecommerce operations. Watch this video to learn about MaxScale, MariaBackup, ColumnStore, ClustrixDB, the MariaDB Platform Managed Service, and associated connectors and data adapters. Learn what each of these products can do for you, and the real-world problems they solve. This video begins by introducing the

database server that forms the core of MariaDB's offering. From there you'll learn about MariaBackup to support backup and recovery, then about MaxScale and Galera Cluster in support of high throughput and scalability. You'll learn about ColumnStore and how it supports analytics processing and column-oriented workloads for business intelligence applications, and about MariaDB's managed service offering in the cloud. Finally, there is the extreme scalability of ClustrixDB and how it enables gaming, social media, IoT, and other extreme workloads. What You Will Learn Get the big picture of the MariaDB family of products Choose the right products with confidence for your application Apply scalability options appropriate to your specific needs Automate backup and recovery through MariaBackup Support business intelligence through columnar storage Unlock extreme scaling using ClustrixDB Who This Video Is For Database users or administrators who are comfortable with MariaDB or MySQL in particular, or with relational databases in general, who want to learn the full range of products available from MariaDB to unlock enterprise-grade performance and scalability.

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 cost-efficient database solutions, scale enterprise operations and reduce overhead business costs with MySQL Key Features Explore the new and advanced features of MySQL 8.0 Use advanced techniques to optimize MySQL performance Create MySQL-based applications for your enterprise with the help of practical examples Book Description Advanced MySQL 8 teaches you to enhance your existing database infrastructure and build various tools to improve your enterprise applications and overall website performance. The book starts with the new and exciting MySQL 8.0 features and how to utilize them for maximum efficiency. As you make your way through the chapters, you will learn to optimize MySQL performance using indexes and advanced data query techniques for large queries. You will also discover MySQL Server 8.0 settings and work with the MySQL data dictionary to boost the performance of your database. In the concluding chapters, you will cover MySQL 8.0 Group Replication, which will enable you to create elastic, highly available, and fault-tolerant replication topologies. You will also

explore backup and recovery techniques for your databases and understand important tips and tricks to help your critical data reach its full potential. By the end of this book, you'll have learned about new MySQL 8.0 security features that allow a database administrator (DBA) to simplify user management and increase the security of their multi-user environments. What you will learn

Explore new and exciting features of MySQL 8.0
Analyze and optimize large MySQL queries
Understand MySQL Server 8.0 settings
Master the deployment of Group Replication and use it in an InnoDB cluster
Monitor large distributed databases
Discover different types of backups and recovery methods for your databases
Explore tips to help your critical data reach its full potential

Who this book is for
Advanced MySQL 8 is for database administrators, data architects, and database developers who want to dive deeper into building advanced database applications in the MySQL environment.

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.

Rust High Performance

Learn to skyrocket the performance of your Rust applications

Learning the MariaDB Ecosystem

Mastering MariaDB

Learn PHP 7

Heading in the Right Direction with MySQL and MariaDB

Although MySQL's source code is open in the sense of being publicly available, it's essentially closed to you if you don't understand it. In this book, Sasha Pachev -- a former member of the MySQL Development Team -- provides a comprehensive tour of MySQL 5 that shows you how to figure out the inner workings of this powerful database. You'll go right to heart of the database to learn how data structures and convenience functions operate, how to add new storage engines and configuration options, and much more. The core of Understanding MySQL Internals begins with an Architecture Overview that provides a brief introduction of how the different components of MySQL work together. You then learn the steps for setting up a working compilable copy of the code that you can change and test at your pleasure. Other sections of the book cover: Core server classes, structures, and API The communication protocol between the client and the server Configuration variables, the controls of the server; includes a tutorial on how to add your own Thread-based request handling -- understanding threads and how they are used in MySQL An overview of MySQL storage engines The storage engine interface for integrating third-party storage engines The table lock manager The parser and optimizer for improving MySQL's performance Integrating a transactional storage engine into MySQL The internals of replication Understanding MySQL Internals provides unprecedented opportunities for developers, DBAs, database application programmers, IT departments, software vendors, and computer science students to learn about the inner workings of this enterprise-proven database. With this book, you will soon reach a new level of comprehension regarding database development that will enable you to accomplish your goals. It's your guide to discovering and improving a great database.

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.

This new book on PHP 7 introduces writing solid, secure, object-oriented code in the new PHP 7: you will create a complete three-tier application using a natural process of building and testing modules within each tier. This practical approach teaches you about app development and introduces PHP features when they are actually needed rather than providing you with abstract theory and contrived examples. In Learn PHP 7, programming examples take advantage of the newest PHP features, including enhanced password encryption using password_hash. This book takes a learn-by-doing approach, providing you with complete coding examples. "Do It" exercises in each chapter provide the opportunity to make adjustments to the example code. The end of chapter programming exercises allow you to develop your own applications using the algorithms demonstrated in the chapter. Each tier is logically and physically separated using object-oriented and dependency injection techniques, thus allowing independent tiers that can be updated with little or no effect on the other tiers. In addition to teaching good programming practices through OOP, there is a strong emphasis on creating secure code. As each chapter is completed, the reader is provide the opportunity to design and create an application reinforcing the concepts learned.

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 Understanding MySQL Internals

Hands-on Data Virtualization with Polybase

OpenVZ Essentials

More than 70 solutions to common Hibernate problems

High-Performance Java Persistence

Step by step guide to monitor, manage, and secure your database engine Key Features Your companion to master all the administration-related tasks in MySQL 8 Ensure high performance and high availability of your MySQL solution using effective replication and backup techniques A comprehensive guide to performing query optimization, security and a whole host of other administrative tasks in MySQL 8 Book Description MySQL is one of the most popular and widely used relational databases in the world today. The recently released version 8.0 brings along some major advancements in the way your MySQL solution can be administered. This handbook will be your companion to understand the newly introduced features in MySQL and how you can leverage them to design a high-performance MySQL solution for your organization. This book starts with a brief introduction to the newly introduced features in MySQL 8, followed by quickly jumping onto the crucial administration topics that you will find useful in your day to day work. Topics such as migrating to MySQL 8, MySQL benchmarking, achieving high performance by implementing the indexing techniques, and optimizing your queries are covered in this book. You will also learn how to perform replication, scale your MySQL solution and implement effective security techniques. A special section on the common and not so common troubleshooting techniques for effective MySQL administration is also covered in this book. By the end of this highly practical book, you

will have all the knowledge you need to tackle any problem you might encounter while administering your MySQL solution. What you will learn Understanding different MySQL 8 data types based on type of contents and storage requirements Best practices for optimal use of features in MySQL 8 Explore globalization configuration and caching techniques to improve performance Create custom storage engine as per system requirements Learn various ways of index implementation for flash memory storages Configure and implement replication along with approaches to use replication as solution Understand how to make your MySQL 8 solution highly available Troubleshoot common issues and identify error codes while using MySQL 8 Who this book is for This book is intended for MySQL administrators who are looking for a handy guide covering all the MySQL administration-related tasks. If you are a DBA looking to get started with MySQL administration, this book will also help you. Knowledge of the basic database concepts is required to get started with this book.

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.

Stuck with bugs, performance problems, crashes, data corruption, and puzzling output? If you're a database programmer or DBA, they're part of your life. The trick is knowing how to quickly recover from them. This unique, example-packed book shows you how to handle an array of vexing problems when working with MySQL. Written by a principal technical support engineer at Oracle, MySQL Troubleshooting provides the background, tools, and expert steps for solving problems from simple to complex—whether data you thought you inserted doesn't turn up in a query, or the entire database is corrupt because of a server failure. With this book in hand, you'll work with more confidence. Understand the source of a problem, even when the solution is simple Handle problems that occur when applications run in multiple threads Debug and fix problems caused by configuration options Discover how operating system tuning can affect your server Use troubleshooting techniques specific to replication issues Get a reference to additional troubleshooting techniques and tools, including third-party solutions Learn best practices for safe and effective troubleshooting—and for preventing problems

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.