
Read PDF MongoDB The Definitive Guide

As recognized, adventure as capably as experience virtually lesson, amusement, as well as accord can be gotten by just checking out a book **MongoDB The Definitive Guide** along with it is not directly done, you could believe even more approaching this life, going on for the world.

We offer you this proper as well as easy pretension to get those all. We offer MongoDB The Definitive Guide and numerous book collections from fictions to scientific research in any way. accompanied by them is this MongoDB The Definitive Guide that can be your partner.

JS705S - JAKOB SCHNEIDER

If you're looking to develop native applications in Kubernetes, this is your guide. Developers and AppOps administrators will learn how to build Kubernetes-native applications that interact directly with the API server to query or update the state of resources. AWS developer advocate Michael Hausenblas and Red Hat principal software engineer Stefan Schimanski explain the characteristics of these apps and show you how to program Kubernetes to build them. You'll explore the basic building blocks of Kubernetes, including the client-go API library and custom resources. All you need to get started is a rudimentary understanding of development and system administration tools and practices, such as package management, the Go programming language, and Git. Walk through Kubernetes API basics and dive into the server's inner structure. Explore Kubernetes's programming interface in Go, including Kubernetes API objects. Learn about custom resources—the central extension tools used in the Kubernetes ecosystem. Use tags to control Kubernetes code generators for custom resources. Write custom controllers and operators and make them production ready. Extend the Kubernetes API surface by implementing a custom API server.

Design, administer, and deploy high-volume and fault-tolerant database applications using MongoDB 4.x. Key Features: Build a powerful and scalable MongoDB database using real industry data. Understand the process of designing NoSQL schema with the latest release of MongoDB 4.x. Explore the ins and outs of MongoDB, including queries, replication, sharding, and vital admin tasks. Book Description: When it comes to managing a high volume of unstructured and non-relational datasets, MongoDB is the defacto database management system (DBMS) for DBAs and data architects. This updated book includes the latest release and covers every feature in MongoDB 4.x, while helping you get hands-on with building a MongoDB database app. You'll get to grips with MongoDB 4.x concepts such as indexes, database design, data modeling, authentication, and

aggregation. As you progress, you'll cover tasks such as performing routine operations when developing a dynamic database-driven website. Using examples, you'll learn how to work with queries and regular database operations. The book will not only guide you through design and implementation, but also help you monitor operations to achieve optimal performance and secure your MongoDB database systems. You'll also be introduced to advanced techniques such as aggregation, map-reduce, complex queries, and generating ad hoc financial reports on the fly. Later, the book shows you how to work with multiple collections as well as embedded arrays and documents, before finally exploring key topics such as replication, sharding, and security using practical examples. By the end of this book, you'll be well-versed with MongoDB 4.x and be able to perform development and administrative tasks associated with this NoSQL database. What you will learn: Understand how to configure and install MongoDB 4.x. Build a database-driven website using MongoDB as the backend. Perform basic database operations and handle complex MongoDB queries. Develop a successful MongoDB database design for large corporate customers with complex requirements. Secure MongoDB database systems by establishing role-based access control with X.509 transport-level security. Optimize reads and writes directed to a replica set or sharded cluster. Perform essential MongoDB administration tasks. Maintain database performance through monitoring. Who this book is for: This book is a MongoDB tutorial for DevOps engineers, database developers, database administrators, system administrators and those who are just getting started with NoSQL and looking to build document-oriented databases and gain real-world experience in managing databases using MongoDB. Basic knowledge of databases and Python is required to get started with this DBMS book.

The definitive guide to building JavaScript-based Web applications from server to browser. Node.js, MongoDB, and AngularJS are three new web development technologies that together provide an easy to im-

plement, fully integrated web development stack. Node.js is a leading server-side programming environment, MongoDB is the most popular NoSQL database, and AngularJS is quickly becoming the leading framework for MVC-based front-end development. Together they allow web programmers to create high-performance sites and applications built completely in JavaScript, from server to client. Node.js, MongoDB and AngularJS Web Development is a complete guide for web programmers who want to integrate these three technologies into full working solutions. It begins with concise, crystal-clear tutorials on each of the three technologies and then quickly moves on to building several common web applications. Readers will learn how to use Node.js and MongoDB to build more scalable, high-performance sites, how to leverage AngularJS's innovative MVC approach to structure more effective pages and applications, and how to use all three together to deliver outstanding next-generation Web solutions.

Get to grips with a new technology, understand what it is and what it can do for you, and then get to work with the most important features and tasks. MongoDB Starter is a fast and practical guide designed to help you start developing high-performance and scalable applications using MongoDB. MongoDB Starter is ideal for developers who are new to MongoDB and who need a no-nonsense guide on how to start working with it. No knowledge of MongoDB is required to follow this book, but some knowledge of C++ would be helpful.

Manage, fine-tune, secure and deploy your MongoDB solution with ease with the help of practical recipes. About This Book: Configure and deploy your MongoDB instance securely, without any hassle. Optimize your database's query performance, perform scale-out operations, and make your database highly available. Practical guide with a recipe-based approach to help you tackle any problem in the application and database administration aspects of MongoDB. Who This Book Is For: Database administrators with a basic understanding of the features of MongoDB and who want to professionally configure,

deploy, and administer a MongoDB database, will find this book essential. If you are a MongoDB developer and want to get into MongoDB administration, this book will also help you. What You Will Learn Install and deploy MongoDB in production Manage and implement optimal indexes Optimize monitoring in MongoDB Fine-tune the performance of your queries Debug and diagnose your database's performance Optimize database backups and recovery and ensure high availability Make your MongoDB instance scalable Implement security and user authentication features in MongoDB Master optimal cloud deployment strategies In Detail MongoDB is a high-performance and feature-rich NoSQL database that forms the backbone of the systems that power many different organizations. Packed with many features that have become essential for many different types of software professional and incredibly easy to use, this cookbook contains more than 100 recipes to address the everyday challenges of working with MongoDB. Starting with database configuration, you will understand the indexing aspects of MongoDB. The book also includes practical recipes on how you can optimize your database query performance, perform diagnostics, and query debugging. You will also learn how to implement the core administration tasks required for high-availability and scalability, achieved through replica sets and sharding, respectively. You will also implement server security concepts such as authentication, user management, role-based access models, and TLS configuration. You will also learn how to back up and recover your database efficiently and monitor server performance. By the end of this book, you will have all the information you need—along with tips, tricks, and best practices—to implement a high-performance MongoDB solution. Style and approach This practical book follows a problem-solution approach to help you tackle any issues encountered while performing MongoDB administrative tasks. Each recipe is detailed, and explained in a very easy to understand manner

Manage your data in a database system designed to support modern application development. The updated edition of this authoritative and accessible guide shows you the many advantages of using document-oriented databases, including how this secure, high-performance system enables flexible data models, high availability, and horizontal scalability. Written by current and former members of the MongoDB team, the third edition is updated for MongoDB 4.0. You'll find substantial updates on querying, indexing, aggregation, replica

sets, ops manager, sharding administration, data administration, durability, monitoring, and security. Authors Shannon Bradshaw (MongoDB) and Kristina Chodorow (Google) provide guidance for database developers, advanced configuration for system administrators, and use cases for a variety of projects. Ideal for NoSQL newcomers and experienced MongoDB users alike, this book also includes many real-world schema design examples. Summary Redis in Action introduces Redis and walks you through examples that demonstrate how to use it effectively. You'll begin by getting Redis set up properly and then exploring the key-value model. Then, you'll dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets. Experienced developers will appreciate chapters on clustering and internal scripting to make Redis easier to use. About the Technology When you need near-real-time access to a fast-moving data stream, key-value stores like Redis are the way to go. Redis expands on the key-value pattern by accepting a wide variety of data types, including hashes, strings, lists, and other structures. It provides lightning-fast operations on in-memory datasets, and also makes it easy to persist to disk on the fly. Plus, it's free and open source. About this book Redis in Action introduces Redis and the key-value model. You'll quickly dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets and discover how to integrate with traditional RDBMS or other NoSQL stores. Experienced developers will appreciate the in-depth chapters on clustering and internal scripting. Written for developers familiar with database concepts. No prior exposure to NoSQL database concepts nor to Redis itself is required. Appropriate for systems administrators comfortable with programming. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside Redis from the ground up Preprocessing real-time data Managing in-memory datasets Pub/sub and configuration Persisting to disk About the Author Dr. Josiah L. Carlson is a seasoned database professional and an active contributor to the Redis community. Table of Contents PART 1 GETTING STARTED Getting to know Redis Anatomy of a Redis web application PART 2 CORE CONCEPTS Commands in Redis Keeping data safe and ensuring performance Using Redis for application support Application components in Redis Search-based applications Building a simple social network PART 3 NEXT

STEPS Reducing memory use Scaling Redis Scripting Redis with Lua Node.js, MongoDB and Angular Web Development The definitive guide to using the MEAN stack to build web applications Node.js is a leading server-side programming environment, MongoDB is the most popular NoSQL database, and Angular is the leading framework for MVC-based front-end development. Together, they provide an easy-to-implement, fully integrated web development stack that allows web programmers to create high-performance sites and applications built completely in JavaScript, from server to client. Updated for Angular 2, Angular 4, and subsequent versions, this new edition of Node.js, MongoDB and Angular Web Development shows you how to integrate these three technologies into complete working solutions. It begins with concise, crystal-clear tutorials on each technology and then quickly moves on to building common web applications. You'll learn how to use Node.js and MongoDB to build more scalable, high-performance sites, how to leverage Angular's innovative MVC approach to structure more effective pages and applications, and how to use all three together to deliver outstanding next-generation Web solutions. Implement a highly scalable and dynamic web server using Node.js and Express Implement a MongoDB data store for your web applications Access and interact with MongoDB from Node.js JavaScript code Learn the basics of TypeScript Define custom Angular directives that extend the HTML language Build server-side web services in JavaScript Implement client-side services that can interact with the Node.js web server Build dynamic browser views that provide rich user interaction Add authenticated user accounts and nested comment components to your web applications and pages Contents at a Glance Part I: Getting Started 1 Introducing the Node.js-to-Angular Stack 2 JavaScript Primer Part II: Learning Node.js 3 Getting Started with Node.js 4 Using Events, Listeners, Timers, and Callbacks in Node.js 5 Handling Data I/O in Node.js 6 Accessing the File System from Node.js 7 Implementing HTTP Services in Node.js 8 Implementing Socket Services in Node.js 9 Scaling Applications Using Multiple Processors in Node.js 10 Using Additional Node.js Modules Part III: Learning MongoDB 11 Understanding NoSQL and MongoDB 12 Getting Started with MongoDB 13 Getting Started with MongoDB and Node.js 14 Manipulating MongoDB Documents from Node.js 15 Accessing MongoDB from Node.js 16 Using Mongoose for Structured Schema and Validation 17 Advanced MongoDB Concepts Part IV: Using Express to Make Life Easier

18 Implementing Express in Node.js 19 Implementing Express Middleware Part V: Learning Angular 20 Jumping into TypeScript 21 Getting Started with Angular 22 Angular Components 23 Expressions 24 Data Binding 25 Built-in Directives Part VI: Advanced Angular 26 Custom Directives 27 Events and Change Detection 28 Implementing Angular Services in Web Applications 29 Creating Your Own Custom Angular Services 30 Having Fun with Angular

Learn how to use, deploy, and maintain Apache Spark with this comprehensive guide, written by the creators of the open-source cluster-computing framework. With an emphasis on improvements and new features in Spark 2.0, authors Bill Chambers and Matei Zaharia break down Spark topics into distinct sections, each with unique goals. You'll explore the basic operations and common functions of Spark's structured APIs, as well as Structured Streaming, a new high-level API for building end-to-end streaming applications. Developers and system administrators will learn the fundamentals of monitoring, tuning, and debugging Spark, and explore machine learning techniques and scenarios for employing MLlib, Spark's scalable machine-learning library. Get a gentle overview of big data and Spark Learn about DataFrames, SQL, and Datasets—Spark's core APIs—through worked examples Dive into Spark's low-level APIs, RDDs, and execution of SQL and DataFrames Understand how Spark runs on a cluster Debug, monitor, and tune Spark clusters and applications Learn the power of Structured Streaming, Spark's stream-processing engine Learn how you can apply MLlib to a variety of problems, including classification or recommendation Learn how to deploy and monitor databases in the cloud, manipulate documents, visualize data, and build applications running on MongoDB using Node.js Key Features Learn the fundamentals of NoSQL databases with MongoDB Create, manage, and optimize a MongoDB database in the cloud using Atlas Use a real-world dataset to gain practical experience of handling big data Book Description MongoDB is one of the most popular database technologies for handling large collections of data. This book will help MongoDB beginners develop the knowledge and skills to create databases and process data efficiently. Unlike other MongoDB books, MongoDB Fundamentals dives into cloud computing from the very start – showing you how to get started with Atlas in the first chapter. You will discover how to modify existing data, add new data into a database, and handle complex queries by creating aggregation

pipelines. As you progress, you'll learn about the MongoDB replication architecture and configure a simple cluster. You will also get to grips with user authentication, as well as techniques for backing up and restoring data. Finally, you'll perform data visualization using MongoDB Charts. You will work on realistic projects that are presented as bitesize exercises and activities, allowing you to challenge yourself in an enjoyable and attainable way. Many of these mini-projects are based around a movie database case study, while the last chapter acts as a final project where you will use MongoDB to solve a real-world problem based on a bike-sharing app. By the end of this book, you'll have the skills and confidence to process large volumes of data and tackle your own projects using MongoDB. What you will learn Set up and use MongoDB Atlas on the cloud Insert, update, delete, and retrieve data from MongoDB Build aggregation pipelines to perform complex queries Optimize queries using indexes Monitor databases and manage user authorization Improve scalability and performance with sharding clusters Replicate clusters, back up your database, and restore data Create data-driven charts and reports from real-time data Who this book is for This book is designed for people who are new to MongoDB. It is suitable for developers, database administrators, system administrators, and cloud architects who are looking to use MongoDB for smooth data processing in the cloud. Although not necessary, basic knowledge of a general programming language and experience with other databases will help you grasp the topics covered more easily.

Imagine what you could do if scalability wasn't a problem. With this hands-on guide, you'll learn how the Cassandra database management system handles hundreds of terabytes of data while remaining highly available across multiple data centers. This expanded second edition—updated for Cassandra 3.0—provides the technical details and practical examples you need to put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's non-relational design, with special attention to data modeling. If you're a developer, DBA, or application architect looking to solve a database scaling issue or future-proof your application, this guide helps you harness Cassandra's speed and flexibility. Understand Cassandra's distributed and decentralized structure Use the Cassandra Query Language (CQL) and cqlsh—the CQL shell Create a working data model and compare it with an equivalent relational model Develop sample applications using

client drivers for languages including Java, Python, and Node.js Explore cluster topology and learn how nodes exchange data Maintain a high level of performance in your cluster Deploy Cassandra on site, in the Cloud, or with Docker Integrate Cassandra with Spark, Hadoop, Elasticsearch, Solr, and Lucene

Master MongoDB - The widely used modern database in a step-by-step, practical, and easy-to-understand approach covering all major topics KEY FEATURES ● In-depth practical demonstration of MongoDB concepts with numerous examples. ● Includes graphical illustrations and visual explanations for MongoDB commands and methods. ● Covers advanced topics such as MongoDB Compass, MongoDB Security, Backup and Restore, and Replication and Sharding. DESCRIPTION MongoDB Complete Guide book starts with the basics of MongoDB, what exactly is MongoDB, and how to use it practically. You will understand how MongoDB is different from the traditional RDBMS. Topics such as installation and configuration of the MongoDB server, MongoDB commands, MongoDB Shell methods, and data types in MongoDB are covered in detail. You will practice how to perform MongoDB CRUD operations, indexing, MongoDB query selectors, projection in MongoDB and projection operators as well as aggregation in a very detailed and step-by-step manner. You learn how to work with MongoDB Compass and some of the advanced MongoDB topics like managing and administering MongoDB, managing the MongoDB process, monitoring and diagnosing MongoDB, backup and restore, MongoDB security, replication and sharding. WHAT YOU WILL LEARN ● Perform write operations, search documents, and define complex queries in MongoDB. ● Perform indexing, aggregation, and data replication. ● End-to-end MongoDB administration along with authentication and authorization. ● Running backups, restoring, and monitoring of MongoDB database enterprise-wide. WHO THIS BOOK IS FOR This book is designed for software developers and server administrators who want to quickly learn MongoDB basics and start applying the knowledge of MongoDB in their business systems. TABLE OF CONTENTS MongoDB Basics 1. Introduction to MongoDB 2. MongoDB Installation and Setup on Windows 3. MongoDB Installation and Setup on Linux (Ubuntu) 4. MongoDB Installation and Setup on macOS 5. Getting started with MongoDB 6. Storage Engines in MongoDB 7. Managing and Administering MongoDB 8. MongoDB Shell Methods 9. Data Types in MongoDB MongoDB Intermediate Level Topics 10. Introduction to MongoDB CRUD Operations 11. MongoDB

Intermediate Concepts 12. Introduction to MongoDB Indexes 13. MongoDB Query Selectors 14. Projection in MongoDB and Projection Operators 15. Aggregation in MongoDB 16. MongoDB Data Manipulations Using MongoDB Compass MongoDB Advanced Level Topics 17. Managing and Administering MongoDB (Advanced Level) 18. Replication in MongoDB 19. Sharding in MongoDB

MongoDB, a cross-platform NoSQL database, is the fastest-growing new database in the world. MongoDB provides a rich document-oriented structure with dynamic queries that you'll recognize from RDBMS offerings such as MySQL. In other words, this is a book about a NoSQL database that does not require the SQL crowd to re-learn how the database world works! MongoDB has reached 1.0 and boasts 50,000+ users. The community is strong and vibrant and MongoDB is improving at a fast rate. With scalable and fast databases becoming critical for today's applications, this book shows you how to install, administer and program MongoDB without pretending SQL never existed.

Leverage the power of MongoDB 4.x to build and administer fault-tolerant database applications Key Features Master the new features and capabilities of MongoDB 4.x Implement advanced data modeling, querying, and administration techniques in MongoDB Includes rich case-studies and best practices followed by expert MongoDB developers Book Description MongoDB is the best platform for working with non-relational data and is considered to be the smartest tool for organizing data in line with business needs. The recently released MongoDB 4.x supports ACID transactions and makes the technology an asset for enterprises across the IT and fintech sectors. This book provides expertise in advanced and niche areas of managing databases (such as modeling and querying databases) along with various administration techniques in MongoDB, thereby helping you become a successful MongoDB expert. The book helps you understand how the newly added capabilities function with the help of some interesting examples and large datasets. You will dive deeper into niche areas such as high-performance configurations, optimizing SQL statements, configuring large-scale sharded clusters, and many more. You will also master best practices in overcoming database failover, and master recovery and backup procedures for database security. By the end of the book, you will have gained a practical understanding of administering database applications both on premises and on the cloud; you will also be able to scale

database applications across all servers. What you will learn Perform advanced querying techniques such as indexing and expressions Configure, monitor, and maintain a highly scalable MongoDB environment Master replication and data sharding to optimize read/write performance Administer MongoDB-based applications on premises or on the cloud Integrate MongoDB with big data sources to process huge amounts of data Deploy MongoDB on Kubernetes containers Use MongoDB in IoT, mobile, and serverless environments Who this book is for This book is ideal for MongoDB developers and database administrators who wish to become successful MongoDB experts and build scalable and fault-tolerant applications using MongoDB. It will also be useful for database professionals who wish to become certified MongoDB professionals. Some understanding of MongoDB and basic database concepts is required to get the most out of this book.

Summary MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology This document-oriented database was built for high availability, supports rich, dynamic schemas, and lets you easily distribute data across multiple servers. MongoDB 3.0 is flexible, scalable, and very fast, even with big data loads. About the Book MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Lots of examples will help you develop confidence in the crucial area of data modeling. You'll also love the deep explanations of each feature, including replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engi-

neer, Doug Garrett is one of the winners of the MongoDB Innovation Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihalis Tsoukalos. Table of Contents PART 1 GETTING STARTED A database for the modern web MongoDB through the JavaScript shell Writing programs using MongoDB PART 2 APPLICATION DEVELOPMENT IN MONGODB Document-oriented data Constructing queries Aggregation Updates, atomic operations, and deletes PART 3 MONGODB MASTERY Indexing and query optimization Text search WiredTiger and pluggable storage Replication Scaling your system with sharding Deployment and administration

This introductory text shows the advantages of using document-oriented databases and demonstrates how MongoDB is a reliable, high-performance system that allows for horizontal scalability. This updated second edition provides guidance for database developers, advanced configuration for system administrators, and an overview of the concepts and use cases for other people on a project.

The Definitive Guide to MongoDB, Third Edition, is updated for MongoDB 3 and includes all of the latest MongoDB features, including the aggregation framework introduced in version 2.2 and hashed indexes in version 2.4. The Third Edition also now includes Node.js along with Python. MongoDB is the most popular of the "Big Data" NoSQL database technologies, and it's still growing. David Hows from 10gen, along with experienced MongoDB authors Peter Membrey and Eelco Plugge, provide their expertise and experience in teaching you everything you need to know to become a MongoDB pro.

Whether you need full-text search or real-time analytics of structured data—or both—the Elasticsearch distributed search engine is an ideal way to put your data to work. This practical guide not only shows you how to search, analyze, and explore data with Elasticsearch, but also helps you deal with the complexities of human language, geolocation, and relationships. If you're a newcomer to both search and distributed systems, you'll quickly learn how to integrate Elasticsearch into your application. More experienced users will pick up lots of advanced techniques. Throughout the book, you'll follow a problem-based approach to learn why, when, and how to use Elasticsearch features. Understand how Elasticsearch interprets data in your documents Index and query your data to take advantage of search concepts such as relevance and word proximity Handle human

language through the effective use of analyzers and queries Summarize and group data to show overall trends, with aggregations and analytics Use geo-points and geo-shapes—Elasticsearch's approaches to geolocation Model your data to take advantage of Elasticsearch's horizontal scalability Learn how to configure and monitor your cluster in production

Whether you're building a social media site or an internal-use enterprise application, this hands-on guide shows you the connection between MongoDB and the business problems it's designed to solve. You'll learn how to apply MongoDB design patterns to several challenging domains, such as ecommerce, content management, and online gaming. Using Python and JavaScript code examples, you'll discover how MongoDB lets you scale your data model while simplifying the development process. Many businesses launch NoSQL databases without understanding the techniques for using their features most effectively. This book demonstrates the benefits of document embedding, polymorphic schemas, and other MongoDB patterns for tackling specific big data use cases, including: Operational intelligence: Perform real-time analytics of business data Ecommerce: Use MongoDB as a product catalog master or inventory management system Content management: Learn methods for storing content nodes, binary assets, and discussions Online advertising networks: Apply techniques for frequency capping ad impressions, and keyword targeting and bidding Social networking: Learn how to store a complex social graph, modeled after Google+ Online gaming: Provide concurrent access to character and world data for a multiplayer role-playing game

Manage the huMONGOus amount of data collected through your web application with MongoDB. This authoritative introduction—written by a core contributor to the project—shows you the many advantages of using document-oriented databases, and demonstrates how this reliable, high-performance system allows for almost infinite horizontal scalability. This updated second edition provides guidance for database developers, advanced configuration for system administrators, and an overview of the concepts and use cases for other people on your project. Ideal for NoSQL newcomers and experienced MongoDB users alike, this guide provides numerous real-world schema design examples. Get started with MongoDB core concepts and vocabulary Perform basic write operations at different levels of safety and speed Create complex queries, with options for limiting, skipping, and sorting re-

sults Design an application that works well with MongoDB Aggregate data, including counting, finding distinct values, grouping documents, and using MapReduce Gather and interpret statistics about your collections and databases Set up replica sets and automatic failover in MongoDB Use sharding to scale horizontally, and learn how it impacts applications Delve into monitoring, security and authentication, backup/restore, and other administrative tasks

With Early Release ebooks, you get books in their earliest form—the author's raw and unedited content as he or she writes—so you can take advantage of these technologies long before the official release of these titles. You'll also receive updates when significant changes are made, new chapters are available, and the final ebook bundle is released. Manage your data in a database system designed to support modern application development. The updated edition of this authoritative and accessible guide shows you the many advantages of using document-oriented databases, including how this secure, high-performance system enables flexible data models, high availability, and horizontal scalability. Written by current and former members of the MongoDB team, the third edition is updated for MongoDB 3.6. You'll find substantial updates on querying, indexing, aggregation, replica sets, ops manager, sharding administration, data administration, durability, monitoring, and security. Authors Shannon Bradshaw (MongoDB) and Kristina Chodorow (Google) provide guidance for database developers, advanced configuration for system administrators, and use cases for a variety of projects. Ideal for NoSQL newcomers and experienced MongoDB users alike, this book also includes many real-world schema design examples.

Learn how to take full advantage of Apache Kafka, the distributed, publish-subscribe queue for handling real-time data feeds. With this comprehensive book, you will understand how Kafka works and how it is designed. Authors Neha Narkhede, Gwen Shapira, and Todd Palino show you how to deploy production Kafka clusters; secure, tune, and monitor them; write rock-solid applications that use Kafka; and build scalable stream-processing applications. Learn how Kafka compares to other queues, and where it fits in the big data ecosystem. Dive into Kafka's internal design Pick up best practices for developing applications that use Kafka. Understand the best way to deploy Kafka in production monitoring, tuning, and maintenance tasks. Learn how to secure a Kafka cluster. You can choose several data access frame-

works when building Java enterprise applications that work with relational databases. But what about big data? This hands-on introduction shows you how Spring Data makes it relatively easy to build applications across a wide range of new data access technologies such as NoSQL and Hadoop. Through several sample projects, you'll learn how Spring Data provides a consistent programming model that retains NoSQL-specific features and capabilities, and helps you develop Hadoop applications across a wide range of use-cases such as data analysis, event stream processing, and workflow. You'll also discover the features Spring Data adds to Spring's existing JPA and JDBC support for writing RDBMS-based data access layers. Learn about Spring's template helper classes to simplify the use of database-specific functionality Explore Spring Data's repository abstraction and advanced query functionality Use Spring Data with Redis (key/value store), HBase (column-family), MongoDB (document database), and Neo4j (graph database) Discover the GemFire distributed data grid solution Export Spring Data JPA-managed entities to the Web as RESTful web services Simplify the development of HBase applications, using a lightweight object-mapping framework Build example big-data pipelines with Spring Batch and Spring Integration

What would happen if you optimized a data store for the operations application developers actually use? You'd arrive at MongoDB, the reliable document-oriented database. With this concise guide, you'll learn how to build elegant database applications with MongoDB and PHP. Written by the Chief Solutions Architect at 10gen—the company that develops and supports this open source database—this book takes you through MongoDB basics such as queries, read-write operations, and administration, and then dives into MapReduce, sharding, and other advanced topics. Get out of the relational database rut, and take advantage of a high-performing system optimized for operations and scale. Learn step-by-step the tools you need to build PHP applications with MongoDB Perform Create, Read, Update, and Delete (CRUD) operations, and learn how to perform queries to retrieve data Administer your database, and access and manipulate data with the MongoDB Shell Use functions to work with sets, arrays, and multiple documents to perform synchronous, asynchronous, and atomic operations Discover PHP's community tools and libraries, and why they're valuable Work with regular expressions, aggregation, MapReduce, replication, and sharding

This is the first book to cover db4o programming in comprehensive detail. Readers are briefed on all of the topics necessary to begin using it in production environments, including installation and configuration, querying and managing objects, performing transactions, and data replication. Newcomers to the topic aren't forgotten, as early chapters are devoted to object database fundamentals, in addition to technical considerations and migration strategies. Complete with numerous C# and Java examples, readers will be able to follow along with the examples regardless of their chosen language.

For too long, developers have worked on disorganized application projects, where every part seemed to have its own build system, and no common repository existed for information about the state of the project. Now there's help. The long-awaited official documentation to Maven is here. Written by Maven creator Jason Van Zyl and his team at Sonatype, *Maven: The Definitive Guide* clearly explains how this tool can bring order to your software development projects. Maven is largely replacing Ant as the build tool of choice for large open source Java projects because, unlike Ant, Maven is also a project management tool that can run reports, generate a project website, and facilitate communication among members of a working team. To use Maven, everything you need to know is in this guide. The first part demonstrates the tool's capabilities through the development, from ideation to deployment, of several sample applications -- a simple software development project, a simple web application, a multi-module project, and a multi-module enterprise project. The second part offers a complete reference guide that includes: The POM and Project Relationships The Build Lifecycle Plugins Project website generation Advanced site generation Reporting Properties Build Profiles The Maven Repository Team Collaboration Writing Plugins IDEs such as Eclipse, IntelliJ, and NetBeans Using and creating assemblies Developing with Maven Archetypes Several sources for Maven have appeared online for some time, but nothing served as an introduction and comprehensive reference guide to this tool -- until now. *Maven: The Definitive Guide* is the ideal book to help you manage development projects for software, web applications, and enterprise applications. And it comes straight from the source.

How does MongoDB help you manage a humongous amount of data collected through your web application? With this authoritative introduction, you'll learn the many advantages of using document-oriented

databases, and discover why MongoDB is a reliable, high-performance system that allows for almost infinite horizontal scalability. Written by engineers from 10gen, the company that develops and supports this open source database, *MongoDB: The Definitive Guide* provides guidance for database developers, advanced configuration for system administrators, and an overview of the concepts and use cases for other people on your project. Learn how easy it is to handle data as self-contained JSON-style documents, rather than as records in a relational database. Explore ways that document-oriented storage will work for your project Learn how MongoDB's schema-free data model handles documents, collections, and multiple databases Execute basic write operations, and create complex queries to find data with any criteria Use indexes, aggregation tools, and other advanced query techniques Learn about monitoring, security and authentication, backup and repair, and more Set up master-slave and automatic failover replication in MongoDB Use sharding to scale MongoDB horizontally, and learn how it impacts applications Get example applications written in Java, PHP, Python, and Ruby

Walk through the basics of Tornado, the high-performance web server known for its speed, simplicity, and scalability on projects large and small. With this hands-on guide, you'll learn how to use Tornado's acclaimed features by working with several example applications. You also get best practices for using Tornado in the real world. Are you interested in creating a scalable social application, real-time analytics engine, or RESTful API—all with the power and simplicity of Python? This book shows you why Tornado is fantastic choice for writing powerful applications that are simple to create, extend, and deploy. Learn how to use Tornado's lightweight and flexible templating language Extend templates to repurpose headers, footers, layout grids, and other content Use persistent storage like MongoDB to store, serve, and edit dynamic content Explore Tornado's ability to make asynchronous web requests Secure your application against cookie and request vulnerabilities Authenticate with external services, using Tornado's auth module Adopt deployment strategies that help harden your application and increase request throughput

Manage your data with a system designed to support modern application development. Updated for MongoDB 4.2, the third edition of this authoritative and accessible guide shows you the advantages of using document-oriented databases. You'll learn

how this secure, high-performance system enables flexible data models, high availability, and horizontal scalability. Authors Shannon Bradshaw, Eoin Brazil, and Kristina Chodorow provide guidance for database developers, advanced configuration for system administrators, and use cases for a variety of projects. NoSQL newcomers and experienced MongoDB users will find updates on querying, indexing, aggregation, transactions, replica sets, ops management, sharding and data administration, durability, monitoring, and security. In six parts, this book shows you how to: Work with MongoDB, perform write operations, find documents, and create complex queries Index collections, aggregate data, and use transactions for your application Configure a local replica set and learn how replication interacts with your application Set up cluster components and choose a shard key for a variety of applications Explore aspects of application administration and configure authentication and authorization Use stats when monitoring, back up and restore deployments, and use system settings when deploying MongoDB Need a quick and easy to understand introduction to MongoDB and NoSQL databases? *MongoDB Basics*, from *The Definitive Guide to MongoDB, 2E*, shows you how a document-oriented database system differs from a relational database, and how to install and get started using it. You'll also learn MongoDB design basics, including geospatial indexing, how to navigate, view, and query your database, and how to use GridFS with a bit of Python.

"*MongoDB and Python*" is a cookbook-style text to help Python programmers work with MongoDB. It is full of useful, practical recipes for solving real-world problems ranging from how to do fast geo queries for location-based apps to efficiently indexing your user documents for social-graph lookups to how best to integrate MongoDB with the Pyramid Web framework.

Work with all aspects of batch processing in a modern Java environment using a selection of Spring frameworks. This book provides up-to-date examples using the latest configuration techniques based on Java configuration and Spring Boot. *The Definitive Guide to Spring Batch* takes you from the "Hello, World!" of batch processing to complex scenarios demonstrating cloud native techniques for developing batch applications to be run on modern platforms. Finally this book demonstrates how you can use areas of the Spring portfolio beyond just Spring Batch 4 to collaboratively develop mission-critical batch processes. You'll see how a new class of use cases and platforms has evolved to have an impact on batch-processing. Data science and big da-

ta have become prominent in modern IT and the use of batch processing to orchestrate workloads has become commonplace. The Definitive Guide to Spring Batch covers how running finite tasks on cloud infrastructure in a standardized way has changed where batch applications are run. Additionally, you'll discover how Spring Batch 4 takes advantage of Java 9, Spring Framework 5, and the new Spring Boot 2 micro-framework. After reading this book, you'll be able to use Spring Boot to simplify the development of your own Spring projects, as well as take advantage of Spring Cloud Task and Spring Cloud Data Flow for added cloud native functionality. Includes a foreword by Dave Syer, Spring Batch project founder. What You'll Learn Discover what is new in Spring Batch 4 Carry out finite batch processing in the cloud using the Spring Batch project Understand the newest configuration techniques based on Java configuration and Spring Boot using practical examples Master batch processing in complex scenarios including in the cloud Develop batch applications to be run on modern platforms Use areas of the Spring portfolio beyond Spring Batch to develop mission-critical batch processes Who This Book Is For Experienced Java and Spring coders new to the Spring Batch platform. This definitive book will be useful in allowing even experienced Spring Batch users and developers to maximize the Spring Batch tool.

Get the definitive guide on Gatsby, the JavaScript framework for building blazing fast websites and applications. Used by Nike, Costa Coffee, and other companies worldwide, Gatsby is emerging as one of the key technologies in the Jamstack (JavaScript, APIs, and markup) ecosystem. With this comprehensive guide, you'll learn how to architect, build, and deploy Gatsby sites independently or with CMSs, commerce systems, and other data sources. Author Preston So begins by showing you how to set up a Gatsby site from scratch. From there, you'll learn ways to use Gatsby's declarative rendering and GraphQL API, build complex offline-enabled sites, and continuously deploy Gatsby sites on a variety of platforms, including Gatsby Cloud. Discover how Gatsby integrates with many data sources and plugins Set up, configure, and architect Gatsby sites using Gatsby's CLI, React, JSX, and GraphQL with high performance out of the box Build an independent Gatsby site based on Markdown and data- and content-driven Gatsby sites that integrate with CMSs and commerce platforms Deploy Gatsby sites with full CI/CD and test coverage on a variety of platforms, including

Netlify, Vercel, and Gatsby Cloud

The Definitive Guide to MongoDB, Second Edition, is updated for the latest version and includes all of the latest MongoDB features, including the aggregation framework introduced in version 2.2 and hashed indexes in version 2.4. MongoDB is the most popular of the "Big Data" NoSQL database technologies, and it's still growing. David Hows from 10gen, along with experienced MongoDB authors Peter Membre and Eelco Plugge, provide their expertise and experience in teaching you everything you need to know to become a MongoDB pro. What you'll learn Set up MongoDB on all major server platforms, including Windows, Linux, OS X, and cloud platforms like Rackspace, Azure, and Amazon EC2 Work with GridFS and the new aggregation framework Work with your data using non-SQL commands Write applications using either PHP or Python Optimize MongoDB Master MongoDB administration, including replication, replication tagging, and tag-aware sharding Who this book is for Database admins and developers who need to get up to speed on MongoDB and its Big Data, NoSQL approach to dealing with data management. Table of Contents- Part I: MongoDB Basics Ch. 1: Introduction to MongoDB Ch. 2: Installing MongoDB Ch. 3: The Data Model Ch. 4: Working with Data Ch. 5: GridFS Part II: Developing with MongoDB Ch. 6: PHP and MongoDB Ch. 7: Python and MongoDB Ch. 8: Advanced Queries Part III: Advanced MongoDB with Big Data Ch. 9: Database Administration Ch. 10: Optimization Ch. 11: Replication Ch. 12: Sharding

Legend has it that Google deploys over two billion application containers a week. How's that possible? Google revealed the secret through a project called Kubernetes, an open source cluster orchestrator (based on its internal Borg system) that radically simplifies the task of building, deploying, and maintaining scalable distributed systems in the cloud. This practical guide shows you how Kubernetes and container technology can help you achieve new levels of velocity, agility, reliability, and efficiency. Authors Kelsey Hightower, Brendan Burns, and Joe Beda—who've worked on Kubernetes at Google and other organizations—explain how this system fits into the lifecycle of a distributed application. You will learn how to use tools and APIs to automate scalable distributed systems, whether it is for online services, machine-learning applications, or a cluster of Raspberry Pi computers. Explore the distributed system challenges that Kubernetes addresses Dive into containerized application development, using containers such as Docker Create and run containers

on Kubernetes, using the docker image format and container runtime Explore specialized objects essential for running applications in production Reliably roll out new software versions without downtime or errors Get examples of how to develop and deploy real-world applications in Kubernetes

Perform fast interactive analytics against different data sources using the Trino high-performance distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Trino. Initially developed by Facebook, open source Trino is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber, and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Trino query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Trino's use cases and learn about tools that will help you connect to Trino and query data Go deeper: Learn Trino's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Trino in production: Secure Trino, monitor workloads, tune queries, and connect more applications; learn how other organizations apply Trino

Outside of the world of enterprise computing, there is one database that enables a huge range of software and hardware to flex relational database capabilities, without the baggage and cost of traditional database management systems. That database is SQLite—an embeddable database with an amazingly small footprint, yet able to handle databases of enormous size. SQLite comes equipped with an array of powerful features available through a host of programming and development environments. It is supported by languages such as C, Java, Perl, PHP, Python, Ruby, TCL, and more. The Definitive Guide to SQLite, Second Edition is devoted to complete coverage of the latest version of this powerful database. It offers a thorough overview of SQLite's capabilities and APIs. The book also uses SQLite as the basis for helping newcomers make their first foray into database development. In only a short time you can be writing programs as diverse as a server-side browser plug-in or the next great iPhone or Android application! Learn about SQLite extensions for C, Java, Perl, PHP, Python, Ru-

by, and Tcl. Get solid coverage of SQLite internals. Explore developing iOS (iPhone) and Android applications with SQLite. SQLite is the solution chosen for thousands of products around the world, from mobile phones and GPS devices to set-top boxes and web browsers. You almost certainly use SQLite every day without even realizing it!

The *Definitive Guide to MongoDB*, Third Edition, is updated for MongoDB 3 and includes all of the latest MongoDB features, including the aggregation framework introduced in version 2.2 and hashed indexes in version 2.4. The Third Edition also now includes Python. MongoDB is the most popular of the "Big Data" NoSQL database technologies, and it's still growing. David Hows from 10gen, along with experienced MongoDB authors Peter Membrey and Eelco Plugge, provide their expertise and experience in teaching you everything you need to know to become a MongoDB pro. Congratulations! You completed the MongoDB application within the given tight timeframe and there is a party to celebrate your application's release into production. Although people are congratulating you at the celebration, you are feeling some uneasiness inside. To complete the project on time required making a lot of assumptions about the data, such as what terms meant and how calculations are derived. In addition, the poor documentation about the application will be of limited use to the support team, and not investigating all of the inherent rules in the data may eventually lead to poorly-performing structures in the not-so-distant future. Now, what if you had a time machine and could

go back and read this book. You would learn that even NoSQL databases like MongoDB require some level of data modeling. Data modeling is the process of learning about the data, and regardless of technology, this process must be performed for a successful application. You would learn the value of conceptual, logical, and physical data modeling and how each stage increases our knowledge of the data and reduces assumptions and poor design decisions. Read this book to learn how to do data modeling for MongoDB applications, and accomplish these five objectives: Understand how data modeling contributes to the process of learning about the data, and is, therefore, a required technique, even when the resulting database is not relational. That is, NoSQL does not mean NoData-Modeling! Know how NoSQL databases differ from traditional relational databases, and where MongoDB fits. Explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts, and learn the basics of adding, querying, updating, and deleting data in MongoDB. Practice a streamlined, template-driven approach to performing conceptual, logical, and physical data modeling. Recognize that data modeling does not always have to lead to traditional data models! Distinguish top-down from bottom-up development approaches and complete a top-down case study which ties all of the modeling techniques together. This book is written for anyone who is working with, or will be working with MongoDB, including business analysts, data modelers,

database administrators, developers, project managers, and data scientists. There are three sections: In Section I, Getting Started, we will reveal the power of data modeling and the tight connections to data models that exist when designing any type of database (Chapter 1), compare NoSQL with traditional relational databases and where MongoDB fits (Chapter 2), explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts (Chapter 3), and explain the basics of adding, querying, updating, and deleting data in MongoDB (Chapter 4). In Section II, Levels of Granularity, we cover Conceptual Data Modeling (Chapter 5), Logical Data Modeling (Chapter 6), and Physical Data Modeling (Chapter 7). Notice the "ing" at the end of each of these chapters. We focus on the process of building each of these models, which is where we gain essential business knowledge. In Section III, Case Study, we will explain both top down and bottom up development approaches and go through a top down case study where we start with business requirements and end with the MongoDB database. This case study will tie together all of the techniques in the previous seven chapters. Nike Senior Data Architect Ryan Smith wrote the foreword. Key points are included at the end of each chapter as a way to reinforce concepts. In addition, this book is loaded with hands-on exercises, along with their answers provided in Appendix A. Appendix B contains all of the book's references and Appendix C contains a glossary of the terms used throughout the text.