

Download Ebook Cassandra Definitive Guide Distributed Scale Read Pdf Free

[Cassandra: The Definitive Guide](#) [Cassandra: The Definitive Guide](#) [Elasticsearch](#) [Elasticsearch: The Definitive Guide](#) **Hadoop: The Definitive Guide** [CockroachDB](#) [Presto: The Definitive Guide](#) **Kafka: The Definitive Guide** [Cassandra: the Definitive Guide, \(Revised\) Third Edition](#) **Elasticsearch: The Definitive Guide** **Mercurial: The Definitive Guide** **HBase** [CouchDB: The Definitive Guide](#) [Cassandra: The Definitive Guide, \(Revised\) Third Edition](#) **Cloud Foundry: The Definitive Guide** **Jenkins: The Definitive Guide** [The Definitive Guide to Jython](#) **Hadoop: The Definitive Guide** **MongoDB: The Definitive Guide** **Spark: The Definitive Guide** **Mercurial: The Definitive Guide** [Security Engineering Managing Channels of Distribution](#) **Asterisk: The Definitive Guide** **Hadoop: The Definitive Guide** [Apache Oracle Design: The Definitive Guide](#) [Tomcat: The Definitive Guide](#) [BEEP](#) [The Definitive Guide to DAX](#) **A Guide to Distribution Theory and Fourier Transforms** [Remote Delivery](#) [Cassandra High Performance Cookbook](#) **Ant** **HADOOP MP3** [Trino: The Definitive Guide](#) **The Definitive Guide to SQLite** **Global Software and IT 802.11 Wireless Networks**

[Cassandra: the Definitive Guide, \(Revised\) Third Edition](#) Jun 22 2022 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 revised third edition--updated for Cassandra 4.0 and new developments in the Cassandra ecosystem, including deployments in Kubernetes with K8ssandra--provides technical details and practical examples to help you put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's nonrelational design, with special attention to data modeling. Developers, DBAs, and application architects looking to solve a database scaling issue or future-proof an application will learn how to 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 Design and develop applications using client drivers Explore cluster

topology and learn how nodes exchange data Maintain a high level of performance in your cluster Deploy Cassandra onsite, in the cloud, or with Docker and Kubernetes Integrate Cassandra with Spark, Kafka, Elasticsearch, Solr, and Lucene **Kafka: The Definitive Guide** Jul 23 2022 Every enterprise application creates data, whether it's log messages, metrics, user activity, outgoing messages, or something else. And how to move all of this data becomes nearly as important as the data itself. If you're an application architect, developer, or production engineer new to Apache Kafka, this practical guide shows you how to use this open source streaming platform to handle real-time data feeds. Engineers from Confluent and LinkedIn who are responsible for developing Kafka explain how to deploy production Kafka clusters, write reliable event-driven microservices, and build scalable stream-processing applications with this platform. Through detailed examples, you'll learn Kafka's design principles, reliability guarantees, key APIs, and architecture details, including the replication protocol, the controller, and the storage layer. Understand publish-subscribe

messaging and how it fits in the big data ecosystem. Explore Kafka producers and consumers for writing and reading messages Understand Kafka patterns and use-case requirements to ensure reliable data delivery Get best practices for building data pipelines and applications with Kafka Manage Kafka in production, and learn to perform monitoring, tuning, and maintenance tasks Learn the most critical metrics among Kafka's operational measurements Explore how Kafka's stream delivery capabilities make it a perfect source for stream processing systems

The Definitive Guide to SQLite Dec 24 2019

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, Ruby, 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!

Hadoop: The Definitive Guide Oct 26 2022

Ready to unlock the power of your data? With this comprehensive guide, you'll learn how to build and maintain reliable, scalable, distributed systems with Apache Hadoop. This book is ideal for programmers looking to analyze datasets of any size, and for administrators who want to set up and run Hadoop clusters. You'll find illuminating case studies that demonstrate how Hadoop is used to solve specific problems. This third edition covers recent changes to Hadoop, including material on the new MapReduce API, as well as MapReduce 2 and its more flexible execution model (YARN). Store large datasets with the Hadoop Distributed File System (HDFS) Run distributed computations with MapReduce Use Hadoop's data and I/O building blocks for compression, data integrity, serialization (including Avro), and persistence Discover common pitfalls and advanced features for

writing real-world MapReduce programs Design, build, and administer a dedicated Hadoop cluster—or run Hadoop in the cloud Load data from relational databases into HDFS, using Sqoop Perform large-scale data processing with the Pig query language Analyze datasets with Hive, Hadoop's data warehousing system Take advantage of HBase for structured and semi-structured data, and ZooKeeper for building distributed systems

Security Engineering May 09 2021 Now that there's software in everything, how can you make anything secure? Understand how to engineer dependable systems with this newly updated classic In *Security Engineering: A Guide to Building Dependable Distributed Systems, Third Edition* Cambridge University professor Ross Anderson updates his classic textbook and teaches readers how to design, implement, and test systems to withstand both error and attack. This book became a best-seller in 2001 and helped establish the discipline of security engineering. By the second edition in 2008, underground dark markets had let the bad guys specialize and scale up; attacks were increasingly on users rather than on technology. The book repeated its success by showing how security engineers can focus on usability. Now the third edition brings it up to date for 2020. As people now go online from phones more than laptops, most servers are in the cloud, online advertising drives the Internet and social networks have taken over much human interaction, many

patterns of crime and abuse are the same, but the methods have evolved. Ross Anderson explores what security engineering means in 2020, including: How the basic elements of cryptography, protocols, and access control translate to the new world of phones, cloud services, social media and the Internet of Things Who the attackers are - from nation states and business competitors through criminal gangs to stalkers and playground bullies What they do - from phishing and carding through SIM swapping and software exploits to DDoS and fake news Security psychology, from privacy through ease-of-use to deception The economics of security and dependability - why companies build vulnerable systems and governments look the other way How dozens of industries went online - well or badly How to manage security and safety engineering in a world of agile development - from reliability engineering to DevSecOps The third edition of Security Engineering ends with a grand challenge: sustainable security. As we build ever more software and connectivity into safety-critical durable goods like cars and medical devices, how do we design systems we can maintain and defend for decades? Or will everything in the world need monthly software upgrades, and become unsafe once they stop?

Hadoop: The Definitive Guide Sep 13 2021 Counsels programmers and administrators for big and small organizations on how to work with large-scale application datasets using Apache Hadoop, discussing its capacity for

storing and processing large amounts of data while demonstrating best practices for building reliable and scalable distributed systems.

Tomcat: The Definitive Guide Nov 03 2020 Jakarta Tomcat is not only the most commonly used open source servlet engine today, it's become the de facto standard by which other servlet engines are measured. Powerful and flexible, it can be used as a stand-alone web server or in conjunction with another server, like Apache or IIS, to run servlets or JSPs. But mastery of Tomcat is not easy: because it's as complex as it is complete. **Tomcat: The Definitive Guide** answers vexing questions that users, administrators, and developers alike have been asking. This concise guide provides much needed information to help harness Tomcat's power and wealth of features. **Tomcat: The Definitive Guide** offers something for everyone who uses Tomcat. System and network administrators will find detailed instructions on installation, configuration, and maintenance. For users, it supplies insightful information on how to deploy Tomcat. And seasoned enterprise Java developers will have a complete reference to setting up, running, and using this powerful software. The book begins with an introduction to the Tomcat server and includes an overview of the three types of server configurations: stand-alone, in-process, and out-of-process. The authors show how directories are laid out, cover the initial setup, and describe how to set the environment variables and modify the configuration files,

concluding with common errors, problems, and solutions. In subsequent chapters, they cover: The server.xml configuration file Java Security manager Authentication schemes and Tomcat users The Secure Socket Layer (SSL) Tomcat JDBC Realms Installing servlets and Java Server Pages Integrating Tomcat with Apache Advanced Tomcat configuration and much more. **Tomcat: The Definitive Guide** covers all major platforms, including Windows, Solaris, Linux, and Mac OS X, contains details on Tomcat configuration files, and has a quick-start guide to get developers up and running with Java servlets and JavaServer Pages. If you've struggled with this powerful yet demanding technology in the past, this book will provide the answers you need.

Ant Apr 27 2020 In 1998 one programmer changed the world of Java. Frustrated by his efforts to create a cross-platform build of Tomcat using the build tools of the day (GNU Make, batch files, and shell scripts), James Duncan Davidson threw together his own build utility on an airplane flight from Europe to the U.S. Named Ant because it was a little thing that could build big things, James's quick-and-dirty solution to his own problem of creating a cross-platform build has evolved into what is perhaps the most widely used build management tool in Java environments.

HBase Mar 19 2022 "HBase: The Definitive Guide" provides the details for evaluating this high-performance, non-relational database, or putting it into practice right away. HBase's

adoption rate is beginning to climb, and IT executives are asking pointed questions about this high-capacity database.

Asterisk: The Definitive Guide Mar 07 2021

Design a complete Voice over IP (VoIP) or traditional PBX system with Asterisk, even if you have only basic telecommunications knowledge. This bestselling guide makes it easy, with a detailed roadmap that shows you how to install and configure this open source software, whether you're upgrading your existing phone system or starting from scratch. Ideal for Linux administrators, developers, and power users, this updated edition shows you how to write a basic dialplan step-by-step, and brings you up to speed on the features in Asterisk 11, the latest long-term support release from Digium. You'll quickly gain working knowledge to build a simple yet inclusive system. Integrate Asterisk with analog, VoIP, and digital telephony systems. Build an interactive dialplan, using best practices for more advanced features. Delve into voicemail options, such as storing messages in a database. Connect to external services including Google Talk, XMPP, and calendars. Incorporate Asterisk features and functions into a relational database to facilitate information sharing. Learn how to use Asterisk's security, call routing, and faxing features. Monitor and control your system with the Asterisk Manager Interface (AMI). Plan for expansion by learning tools for building distributed systems.

Elasticsearch: The Definitive Guide May 21

2022 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.

Jenkins: The Definitive Guide Nov 15 2021 Streamline software development with Jenkins, the popular Java-based open source tool that has revolutionized the way teams think about Continuous Integration (CI). This complete guide shows you how to automate your build,

integration, release, and deployment processes with Jenkins—and demonstrates how CI can save you time, money, and many headaches. Ideal for developers, software architects, and project managers, *Jenkins: The Definitive Guide* is both a CI tutorial and a comprehensive Jenkins reference. Through its wealth of best practices and real-world tips, you'll discover how easy it is to set up a CI service with Jenkins. Learn how to install, configure, and secure your Jenkins server. Organize and monitor general-purpose build jobs. Integrate automated tests to verify builds, and set up code quality reporting. Establish effective team notification strategies and techniques. Configure build pipelines, parameterized jobs, matrix builds, and other advanced jobs. Manage a farm of Jenkins servers to run distributed builds. Implement automated deployment and continuous delivery.

Elasticsearch Dec 28 2022 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

Hadoop: The Definitive Guide Feb 06 2021 Discover how Apache Hadoop can unleash the power of your data. This comprehensive resource shows you how to build and maintain reliable, scalable, distributed systems with the Hadoop framework -- an open source implementation of MapReduce, the algorithm on which Google built its empire. Programmers will find details for analyzing datasets of any size, and administrators will learn how to set up and run Hadoop clusters. This revised edition covers recent changes to Hadoop, including new features such as Hive, Sqoop, and Avro. It also provides illuminating case studies that illustrate how Hadoop is used to solve specific problems. Looking to get the most out of your data? This is your book. Use the Hadoop Distributed File System (HDFS) for storing large datasets, then run distributed computations over those datasets with

MapReduce Become familiar with Hadoop's data and I/O building blocks for compression, data integrity, serialization, and persistence Discover common pitfalls and advanced features for writing real-world MapReduce programs Design, build, and administer a dedicated Hadoop cluster, or run Hadoop in the cloud Use Pig, a high-level query language for large-scale data processing Analyze datasets with Hive, Hadoop's data warehousing system Take advantage of HBase, Hadoop's database for structured and semi-structured data Learn ZooKeeper, a toolkit of coordination primitives for building distributed systems "Now you have the opportunity to learn about Hadoop from a master -- not only of the technology, but also of common sense and plain talk." --Doug Cutting, Cloudera

The Definitive Guide to Jython Oct 14 2021 Jython is an open source implementation of the high-level, dynamic, object-oriented scripting language Python seamlessly integrated with the Java platform. The predecessor to Jython, JPython, is certified as 100% Pure Java. Jython is freely available for both commercial and noncommercial use and is distributed with source code. Jython is complementary to Java. The Definitive Guide to Jython, written by the official Jython team leads, covers Jython 2.5 (or 2.5.x)—from the basics to more advanced features. This book begins with a brief introduction to the language and then journeys through Jython's different features and uses. The Definitive Guide to Jython is organized for

beginners as well as advanced users of the language. The book provides a general overview of the Jython language itself, but it also includes intermediate and advanced topics regarding database, web, and graphical user interface (GUI) applications; Web services/SOA; and integration, concurrency, and parallelism, to name a few.

Cassandra: The Definitive Guide Mar 02 2023 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

MongoDB: The Definitive Guide Aug 12 2021 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 results 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

Global Software and IT Nov 22 2019 Based on the author's first-hand experience and expertise, this book offers a proven framework for global software engineering. Readers will learn best practices for managing a variety of software projects, coordinating the activities of several locations across the globe while accounting for cultural differences. Most importantly, readers will learn how to engineer a first-rate software product as efficiently as possible by fully leveraging global personnel and resources. *Global Software and IT* takes a unique approach that works for projects of any size, examining such critical topics as: Executing a seamless project across multiple locations Mitigating the risks of off-shoring Developing and implementing processes for global development Establishing practical outsourcing guidelines Fostering effective collaboration and communication across continents and culture This book provides a balanced framework for planning global development, covering topics such as managing people in distributed sites and managing a project across locations. It delivers a comprehensive business model that is beneficial to anyone looking for the most cost-effective, efficient way to engineer good software products.

Cloud Foundry: The Definitive Guide Dec 16 2021 How can Cloud Foundry help you develop and deploy business-critical applications and

tasks with velocity? This practical guide demonstrates how this open source, cloud-native application platform not only significantly reduces the develop-to-deploy cycle time, but also raises the value line for application operators by changing the way applications and supporting services are deployed and run. Learn how Cloud Foundry can help you improve your product velocity by handling many of the essential tasks required to run applications in production. Author Duncan Winn shows DevOps and operations teams how to configure and run Cloud Foundry at scale. You'll examine Cloud Foundry's technical concepts—including how various platform components interrelate—and learn how to choose your underlying infrastructure, define the networking architecture, and establish resiliency requirements. This book covers: Cloud-native concepts that make the app build, test, deploy, and scale faster How to deploy Cloud Foundry and the BOSH release engineering toolchain Concepts and components of Cloud Foundry's runtime architecture Cloud Foundry's routing mechanisms and capabilities The platform's approach to container tooling and orchestration BOSH concepts, deployments, components, and commands Basic tools and techniques for debugging the platform Recent and soon-to-emerge features of Cloud Foundry

Remote Delivery Jun 29 2020 This book records the author's years of experience in the software industry. In his own practices, the author has

found that the distributed work pattern has become increasingly popular in more and more work environments, either between vendors and customers or between different teams inside a company. This means that all practitioners in the software industry need to adapt to this new way of communication and collaboration and get skilled enough to meet the greater challenges in integrating the distributed work pattern with agile software delivery. By centering on the difficulties in communication and collaboration between distributed teams, this book digs into the reasons why so many remote delivery projects end up anticlimactic and provides solutions for readers' reference. It also cites successful cases in promoting agile development in distributed teams, which has been a vexing problem for many software development companies. In addition, readers can find suggestions and measures for building self-managing teams in this book. Remote Delivery: A Guide to Software Delivery through Collaboration between Distributed Teams is a very practical guide for software delivery teams with their members distributed in different places and companies engaged in software customization. Developers, QAs, product managers, and project leaders can also be inspired by this book.

Presto: The Definitive Guide Aug 24 2022 Perform fast interactive analytics against different data sources using the Presto 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 Presto. Initially developed by Facebook, open source Presto 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 Presto query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Presto's use cases and learn about tools that will help you connect to Presto and query data Go deeper: Learn Presto's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Presto in production: Secure Presto, monitor workloads, tune queries, and connect more applications; learn how other organizations apply Presto

MP3 Feb 24 2020 You'll find a detailed examination of the codec itself and its development, as well as how MP3 files compare with other common compression formats."

Mercurial: The Definitive Guide Apr 20 2022 This instructive book takes you step by step through ways to track, merge, and manage both open source and commercial software projects with Mercurial, using Windows, Mac OS X, Linux, Solaris, and other systems. Mercurial is

the easiest system to learn when it comes to distributed revision control. And it's a very flexible tool that's ideal whether you're a lone programmer working on a small project, or part of a huge team dealing with thousands of files. Mercurial permits a countless variety of development and collaboration methods, and this book offers several concrete suggestions to get you started. This guide will help you: Learn the basics of working with a repository, changesets, and revisions Merge changes from separate repositories Set up Mercurial to work with files on a daily basis, including which ones to track Get examples and tools for setting up various workflow models Manage a project that's making progress on multiple fronts at once Find and fix mistakes by isolating problem sources Use hooks to perform actions automatically in response to repository events Customize the output of Mercurial Mercurial: The Definitive Guide maintains a strong focus on simplicity to help you learn Mercurial quickly and thoroughly.

Cassandra: The Definitive Guide, (Revised) Third Edition Jan 17 2022 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 revised third edition--updated for Cassandra 4.0 and new developments in the Cassandra ecosystem, including deployments in Kubernetes with K8ssandra--provides technical

details and practical examples to help you put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's nonrelational design, with special attention to data modeling. Developers, DBAs, and application architects looking to solve a database scaling issue or future-proof an application will learn how to 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 Design and develop applications using client drivers Explore cluster topology and learn how nodes exchange data Maintain a high level of performance in your cluster Deploy Cassandra onsite, in the cloud, or with Docker and Kubernetes Integrate Cassandra with Spark, Kafka, Elasticsearch, Solr, and Lucene

Oracle Design: The Definitive Guide Dec 04 2020 This book focuses exclusively on Oracle database design. It covers the most up-to-date Oracle issues and technologies, including massively parallel processors, very large databases, data warehouses, client-server, and distributed database. The design advice is detailed and thorough. The book delves deeply into design issues and gives advice that will have a major impact on your database and system performance.

CouchDB: The Definitive Guide Feb 18 2022

Three of CouchDB's creators show you how to use this document-oriented database as a standalone application framework or with high-volume, distributed applications. With its simple model for storing, processing, and accessing data, CouchDB is ideal for web applications that handle huge amounts of loosely structured data. That alone would stretch the limits of a relational database, yet CouchDB offers an open source solution that's reliable, scales easily, and responds quickly. CouchDB works with self-contained data that has loose or ad-hoc connections. It's a model that fits many real-world items, such as contacts, invoices, and receipts, but you'll discover that this database can easily handle data of any kind. With this book, you'll learn how to work with CouchDB through its RESTful web interface, and become familiar with key features such as simple document CRUD (create, read, update, delete), advanced MapReduce, deployment tuning, and more. Understand the basics of document-oriented storage and manipulation Interact with CouchDB entirely through HTTP using its RESTful interface Model data as self-contained JSON documents Handle evolving data schemas naturally Query and aggregate data in CouchDB using MapReduce views Replicate data between nodes Tune CouchDB for increased performance and reliability **Mercurial: The Definitive Guide** Jun 10 2021 This instructive book takes you step by step through ways to track, merge, and manage both

open source and commercial software projects with Mercurial, using Windows, Mac OS X, Linux, Solaris, and other systems. Mercurial is the easiest system to learn when it comes to distributed revision control. And it's a very flexible tool that's ideal whether you're a lone programmer working on a small project, or part of a huge team dealing with thousands of files. Mercurial permits a countless variety of development and collaboration methods, and this book offers several concrete suggestions to get you started. This guide will help you: Learn the basics of working with a repository, changesets, and revisions Merge changes from separate repositories Set up Mercurial to work with files on a daily basis, including which ones to track Get examples and tools for setting up various workflow models Manage a project that's making progress on multiple fronts at once Find and fix mistakes by isolating problem sources Use hooks to perform actions automatically in response to repository events Customize the output of Mercurial Mercurial: The Definitive Guide maintains a strong focus on simplicity to help you learn Mercurial quickly and thoroughly. *The Definitive Guide to DAX* Sep 01 2020 This comprehensive and authoritative guide will teach you the DAX language for business intelligence, data modeling, and analytics. Leading Microsoft BI consultants Marco Russo and Alberto Ferrari help you master everything from table functions through advanced code and model optimization. You'll learn exactly

what happens under the hood when you run a DAX expression, how DAX behaves differently from other languages, and how to use this knowledge to write fast, robust code. If you want to leverage all of DAX's remarkable power and flexibility, this no-compromise "deep dive" is exactly what you need. Perform powerful data analysis with DAX for Microsoft SQL Server Analysis Services, Excel, and Power BI Master core DAX concepts, including calculated columns, measures, and error handling Understand evaluation contexts and the CALCULATE and CALCULATETABLE functions Perform time-based calculations: YTD, MTD, previous year, working days, and more Work with expanded tables, complex functions, and elaborate DAX expressions Perform calculations over hierarchies, including parent/child hierarchies Use DAX to express diverse and unusual relationships Measure DAX query performance with SQL Server Profiler and DAX Studio

Trino: The Definitive Guide Jan 25 2020

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 *Elasticsearch: The Definitive Guide* Nov 27 2022 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

Managing Channels of Distribution Apr 08 2021

"Channels of distribution is one of the hottest areas in marketing and sales today. And no one understands the subject better than Ken Rolnicki! *Managing Channels of Distribution* supplies a much-needed source of knowledge and expertise that professionals can rely on. Based on case studies and real-life experience, the book explains the complexities of managing multiple channels -- distributors, dealers, manufacturer's reps, VARs, private labels, brokers, wholesalers, retailers, and all the rest. In the process, Rolnicki explores both macro and micro business influences that affect channel effectiveness. Special attention is paid to the frustrating areas of channel power and conflict, the dangerous issue of legalities, and the most critical topic of all -- the channel design sequence."

BEEP Oct 02 2020 A developer's handbook for understanding the BEEP framework and how to put it to best use. The API's for C, Java, and Tel are explored.

Spark: The Definitive Guide Jul 11 2021

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

CockroachDB Sep 25 2022 Get the lowdown on CockroachDB, the elastic SQL database built to handle the demands of today's data-driven world. With this practical guide, software developers, architects, and DevOps teams will discover the advantages of building on a

distributed SQL database. You'll learn how to create applications that scale elastically and provide seamless delivery for end users while remaining exceptionally resilient and indestructible. Written from scratch for the cloud and architected to scale elastically to handle the demands of cloud native and open source, CockroachDB makes it easier to build and scale modern applications. If you're familiar with distributed systems, you'll quickly discover the benefits of strong data correctness and consistency guarantees as well as optimizations for delivering ultralow latencies to globally distributed end users. With this thorough guide, you'll learn how to: Plan and build applications for distributed infrastructure, including data modeling and schema design Migrate data into CockroachDB Read and write data and run ACID transactions across distributed infrastructure Optimize queries for performance across geographically distributed replicas Plan a CockroachDB deployment for resiliency across single-region and multiregion clusters Secure, monitor, and optimize your CockroachDB deployment.

802.11 Wireless Networks Oct 22 2019 As we all know by now, wireless networks offer many advantages over fixed (or wired) networks. Foremost on that list is mobility, since going wireless frees you from the tether of an Ethernet cable at a desk. But that's just the tip of the cable-free iceberg. Wireless networks are also more flexible, faster and easier for you to use, and more affordable to deploy and

maintain. The de facto standard for wireless networking is the 802.11 protocol, which includes Wi-Fi (the wireless standard known as 802.11b) and its faster cousin, 802.11g. With easy-to-install 802.11 network hardware available everywhere you turn, the choice seems simple, and many people dive into wireless computing with less thought and planning than they'd give to a wired network. But it's wise to be familiar with both the capabilities and risks associated with the 802.11 protocols. And *802.11 Wireless Networks: The Definitive Guide, 2nd Edition* is the perfect place to start. This updated edition covers everything you'll ever need to know about wireless technology. Designed with the system administrator or serious home user in mind, it's a no-nonsense guide for setting up 802.11 on Windows and Linux. Among the wide range of topics covered are discussions on: deployment considerations network monitoring and performance tuning wireless security issues how to use and select access points network monitoring essentials wireless card configuration security issues unique to wireless networks With wireless technology, the advantages to its users are indeed plentiful. Companies no longer have to deal with the hassle and expense of wiring buildings, and households with several computers can avoid fights over who's online. And now, with *802.11 Wireless Networks: The Definitive Guide, 2nd Edition*, you can integrate wireless technology into your current infrastructure with the utmost

confidence.

Cassandra High Performance Cookbook May 29 2020 Over 150 recipes to design and optimize large scale Apache Cassandra deployments.

Apache Jan 05 2021 Describes the history of the Web server platform and covers downloading and compiling, configuring and running the program on UNIX, writing specialized modules, and establishing security routines.

A Guide to Distribution Theory and Fourier Transforms Jul 31 2020 This important book provides a concise exposition of the basic ideas of the theory of distribution and Fourier transforms and its application to partial differential equations. The author clearly presents the ideas, precise statements of theorems, and explanations of ideas behind the proofs. Methods in which techniques are used in applications are illustrated, and many problems are included. The book also

introduces several significant recent topics, including pseudodifferential operators, wave front sets, wavelets, and quasicrystals.

Background mathematical prerequisites have been kept to a minimum, with only a knowledge of multidimensional calculus and basic complex variables needed to fully understand the concepts in the book. A Guide to Distribution Theory and Fourier Transforms can serve as a textbook for parts of a course on Applied Analysis or Methods of Mathematical Physics, and in fact it is used that way at Cornell.

HADOOP Mar 27 2020

Cassandra: The Definitive Guide Jan 29 2023 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 third

edition—updated for Cassandra 4.0—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 nonrelational 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