

Cassandra The Definitive Guide

Thank you categorically much for downloading Cassandra The Definitive Guide. Most likely you have knowledge that, people have seen numerous periods for their favorite books subsequently this Cassandra The Definitive Guide, but end stirring in harmful downloads.

Rather than enjoying a good book bearing in mind a cup of coffee in the afternoon, then again they juggled next some harmful virus inside their computer. Cassandra The Definitive Guide is friendly in our digital library an online access to it is set as public therefore you can download it instantly. Our digital library saves in multiple countries, allowing you to acquire the most less latency times to download any of our books afterward this one. Merely said, the Cassandra The Definitive Guide is universally compatible taking into account any devices to read.

Cassandra: the Definitive Guide Jeff Carpenter 2020-07-31 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

Ant: The Definitive Guide Steve Holzner 2005-04-13 Soon after its launch, Ant succeeded in taking the Java world by storm, becoming the most widely used tool for building applications in Java environments. Like most popular technologies, Ant quickly went through a series of early revision cycles. With each new version, more functionality was added, and more complexity was introduced. Ant evolved from a simple-to-learn build tool into a full-fledged testing and deployment environment. Ant: The Definitive Guide has been reworked, revised and expanded upon to reflect this evolution. It documents the new ways that Ant is being applied, as well as the array of optional tasks that Ant supports. In fact, this new second edition covers everything about this extraordinary build management tool from downloading and installing, to using Ant to test code. Here are just a few of the features you'll find detailed in this comprehensive, must-have guide: Developing conditional builds, and handling error conditions Automatically retrieving source code from version control systems Using Ant with XML files Using Ant with JavaServer Pages to build Web applications Using Ant with Enterprise JavaBeans to build enterprise applications Far exceeding its predecessor in terms of information and detail, Ant: The Definitive Guide, 2nd Edition is a must-have for Java developers unfamiliar with the latest advancements in Ant technology. With this book at your side, you'll soon be up to speed on the premiere tool for cross-platform development. Author Steve Holzner is an award-winning author who's been writing about Java topics since the language first appeared; his books have sold more than 1.5 million copies worldwide.

M48 Patton Robert Griffin 2020-01-30 A fully illustrated and historically informative guide to building scale models of the U.S. Army's iconic Cold War main battle tank. The M48 Patton main battle tank was one of the longest-serving and most successful designs in the U.S. military. A significant advance from the Sherman and M47, the M48 formed the backbone of American armored forces during the early years of the Cold War. It saw combat in Vietnam, during the Indo-Pakistan wars and in the Middle East, especially during the Iran-Iraq War. Archive photos of the M48 in service and extensively researched color profile illustrations depict the tank throughout its operational life. A large part of the book showcases available model kits and aftermarket products, complemented by a gallery of beautifully constructed and painted models in various scales. Technical details as well as modifications introduced during production and in the field are also examined and provide everything the modeler needs to recreate an accurate representation of the M48.

Programming Hive Edward Capriolo 2012-09-26 Describes the features and functions of Apache Hive, the data infrastructure for Hadoop.

Database Internals Alex Petrov 2019-09-13 When it comes to choosing, using, and maintaining a database, understanding its internals is essential. But with so many distributed databases and tools available today, it's often difficult to understand what each one offers and how they differ. With this practical guide, Alex Petrov guides developers through the concepts behind modern database and storage engine internals. Throughout the book, you'll explore relevant material gleaned from numerous books, papers, blog posts, and the source code of several open source databases. These resources are listed at the end of parts one and two. You'll discover that the most significant distinctions among many modern databases reside in subsystems that determine how storage is organized and how data is distributed. This book examines: Storage engines: Explore storage classification and taxonomy, and dive into B-Tree-based and immutable Log Structured storage engines, with differences and use-cases for each Storage building blocks: Learn how database files are organized to build efficient storage, using auxiliary data structures such as Page Cache, Buffer Pool and Write-Ahead Log Distributed systems: Learn step-by-step how nodes and processes connect and build complex communication patterns Database clusters: Which consistency models are commonly used by modern databases and how distributed storage systems achieve consistency

Kafka: The Definitive Guide Neha Narkhede 2017-08-31 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

Practical Cassandra Russell Bradberry 2013 "Eric and Russell were early adopters of Cassandra at SimpleReach. In Practical Cassandra, you benefit from their experience in the trenches administering Cassandra, developing against it, and building one of the first CQL drivers. If you are deploying Cassandra soon, or you inherited a Cassandra cluster to tend, spend some time with the deployment, performance tuning, and maintenance chapters... If you are new to Cassandra, I highly recommend the chapters on data modeling and CQL." --From the Foreword by Jonathon Ellis, Apache Cassandra Chair Build and Deploy Massively Scalable, Super-fast Data Management Applications with Apache Cassandra Practical Cassandra is the first hands-on developer's guide to building Cassandra systems and applications that deliver breakthrough speed, scalability, reliability, and performance. Fully up to date, it reflects the latest versions of Cassandra--including Cassandra Query Language (CQL), which dramatically lowers the learning curve for Cassandra developers. Pioneering Cassandra developers and Datastax MVPs Russell Bradberry and Eric Lubow walk you through every step of building a real production application that can store enormous amounts of structured, semi-structured, and unstructured data. Drawing on their exceptional expertise, Bradberry and Lubow share practical insights into issues ranging from querying to deployment, management, maintenance, monitoring, and troubleshooting. The authors cover key

issues, from architecture to migration, and guide you through crucial decisions about configuration and data modeling. They provide tested sample code, detailed explanations of how Cassandra works "under the covers," and new case studies from three cutting-edge users: Ooyala, Hailo, and eBay. Coverage includes Understanding Cassandra's approach, architecture, key concepts, and primary use cases— and why it's so blazingly fast Getting Cassandra up and running on single nodes and large clusters Applying the new design patterns, philosophies, and features that make Cassandra such a powerful data store Leveraging CQL to simplify your transition from SQL-based RDBMSes Deploying and provisioning through the cloud or on bare-metal hardware Choosing the right configuration options for each type of workload Tweaking Cassandra to get maximum performance from your hardware, OS, and JVM Mastering Cassandra's essential tools for maintenance and monitoring Efficiently solving the most common problems with Cassandra deployment, operation, and application development

Cassandra: The Definitive Guide Jeff Carpenter 2016-06-29 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

CASSANDRA EBEN HEWITT. JEFF CARPENTER 2016

The Definitive Guide to MongoDB Eelco Plugge 2015-12-11 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.

Jenkins: The Definitive Guide John Ferguson Smart 2011-07-12 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

Data Modeling for MongoDB Steve Hoberman 2014-06-01 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 NoDataModeling! 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.

Cassandra Jeff Carpenter 2016-07-08 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

Cassandra: The Definitive Guide Jeff Carpenter 2016-06-29 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

Cassandra

Jeff Carpenter. Eben Hewitt 2016

The Palmistry Bible Jane Struthers 2005 Discover in this illuminating volume how you can use palm reading for self-exploration and personal growth.

Cassandra: The Definitive Guide Eben Hewitt 2010-11-12 What could you do with data if scalability wasn't a problem? With this hands-on guide, you'll learn how Apache Cassandra handles hundreds of terabytes of data while remaining highly available across multiple data centers -- capabilities that have attracted Facebook, Twitter, and other data-intensive companies. Cassandra: The Definitive Guide provides the technical details and practical examples you need to assess this database management system and put it to work in a production environment. Author Eben Hewitt demonstrates the advantages of Cassandra's nonrelational design, and pays special attention to data modeling. If you're a developer, DBA, application architect, or manager looking to solve a database scaling issue or future-proof your application, this guide shows you how to harness Cassandra's speed and flexibility. Understand the tenets of Cassandra's column-oriented structure Learn how to write, update, and read Cassandra data Discover how to add or remove nodes from the cluster as your application requires Examine a working application that translates from a relational model to Cassandra's data model Use examples for writing clients in Java, Python, and C# Use the JMX interface to monitor a cluster's usage, memory patterns, and more Tune memory settings, data storage, and caching for better performance

OSSEC Host-Based Intrusion Detection Guide Rory Bray 2008-04-09 This book is the definitive guide on the OSSEC Host-based Intrusion Detection system and frankly, to really use OSSEC you are going to need a definitive guide. Documentation has been available since the start of the OSSEC project but, due to time constraints, no formal book has been created to outline the various features and functions of the OSSEC product. This has left very important and powerful features of the product undocumented...until now! The book you are holding will show you how to install and configure OSSEC on the operating system of your choice and provide detailed examples to help prevent and mitigate attacks on your systems. -- Stephen Northcutt OSSEC determines if a host has been compromised in this manner by taking the equivalent of a picture of the host machine in its original, unaltered state. This "picture" captures the most relevant information about that machine's configuration. OSSEC saves this "picture" and then constantly compares it to the current state of that machine to identify anything that may have changed from the original configuration. Now, many of these changes are necessary, harmless, and authorized, such as a system administrator installing a new software upgrade, patch, or application. But, then there are the not-so-harmless changes, like the installation of a rootkit, trojan horse, or virus. Differentiating between the harmless and the not-so-harmless changes determines whether the system administrator or security professional is managing a secure, efficient network or a compromised network which might be funneling credit card numbers out to phishing gangs or storing massive amounts of pornography creating significant liability for that organization. Separating the wheat from the chaff is by no means an easy task. Hence the need for this book. The book is co-authored by Daniel Cid, who is the founder and lead developer of the freely available OSSEC host-based IDS. As such, readers can be certain they are reading the most accurate, timely, and insightful information on OSSEC. * Nominee for Best Book Bejtlich read in 2008! * <http://taosecurity.blogspot.com/2008/12/best-book-bejtlich-read-in-2008.html> • Get Started with OSSEC Get an overview of the features of OSSEC including commonly used terminology, pre-install preparation, and deployment considerations. • Follow Step-by-Step Installation Instructions Walk through the installation process for the "local", "agent", and "server" install types on some of the most popular operating systems available. • Master Configuration Learn the basic configuration options for your install type and learn how to monitor log files, receive remote messages, configure email notification, and configure alert levels. • Work With Rules Extract key information from logs using decoders and how you can leverage rules to alert you of strange occurrences on your network. • Understand System Integrity Check and Rootkit Detection Monitor binary executable files, system configuration files, and the Microsoft Windows registry. • Configure Active Response Configure the active response actions you want and bind the actions to specific rules and sequence of events. • Use the OSSEC Web User Interface Install, configure, and use the community-developed, open source web interface available for OSSEC. • Play in the OSSEC VMware Environment Sandbox • Dig Deep into Data Log Mining Take the "high art of log analysis to the next level by breaking the dependence on the lists of strings or patterns to look for in the logs.

Weight Loss Burton Goldberg 2000 Explains how to use enzymes, aromatherapy, diet, nutrition, Ayurvedic medicine, glandular extracts, exercise, and supplements to help lose weight permanently.

The Trader's Book of Volume: The Definitive Guide to Volume Trading Mark Leibovit 2011-01-07 Learn how to translate the "language" of volume! Mark Leibovit, a leading market strategist and technical analyst with more than 35 years of trading experience, possesses a solid track record of predicting important movements in the financial market—including Black Monday of 1987, the bear markets of 2000 and 2008, and the "flash crash" of May 2010. Now, with The Trader's Book of Volume, his secrets are yours! Focusing exclusively on volume technical analysis, The Trader's Book of Volume describes the basics of volume, explains how to use it to identify and assess the strength of trade-worthy trends, and provides in-depth techniques and strategies for trading volume indicators for profit. With more than 400 charts and graphs, The Trader's Book of Volume also exhaustively illustrates how readers can profit from a wide array of volume indicators, including: Broad Market Volume Indicators—Cumulative Volume Index, ARMS Index, Upside-Downside Volume, Nasdaq/ NYSE Volume Ratio, Yo-Yo Indicator Volume Indicators—Accumulation/ Distribution, Intraday Intensity, Negative Volume Index, On-Balance Volume, Open Interest Volume Oscillators—Klinger Oscillator, Chaikin Money Flow, Ease of Movement, Volume Oscillator Leibovit Volume Reversal Indicator™, the author's proprietary methodology Under the author's expert guidance, you can seamlessly incorporate Volume Analysis into your day-to-day trading program. Without a proper approach to Volume Analysis, Leibovit asserts, you're essentially trading in the "land of the blind." Use The Trader's Book of Volume to gain the clearest view possible of market trends and react to them with the confidence and smarts for consistent trading success—and avoid every market crash the future holds.

The Book of Esports William Collis 2020-08-04 The definitive guide to the modern world of competitive gaming and the official history of Esports™. Almost overnight, esports—or competitive video games—have exploded into the largest entertainment and sporting phenomenon in human history. The Book of Esports answers: What exactly are esports, and how did they become so popular so quickly? Why did blockbuster video games like League of Legends, Fortnite and Starcraft succeed? Where exactly is all this video gaming headed? What do gamers and college students need to know to position themselves for success in the industry? How do you create a billion-dollar esports business? What strategic choices drive success in the modern gaming industry? Can video games really get your kid into college? (All expenses paid, of course...) Whether you are a lifelong gamer, a curious Fortnite parent, or a businessperson seeking to understand the marketing opportunities of this multibillion-dollar phenomenon, The Book of Esports charts the rise of this exciting new industry, for the first time ever crafting a comprehensive overview of esports and its implications for human competition—and even the future of humanity itself. Gaming luminary and Harvard MBA William Collis has painstakingly translated esports' mysteries into a detailed and accessible testament for today. Featuring select interviews from the biggest names in the industry, The Book of Esports weaves tales of trust, betrayal, and superhuman reflexes into predictive frameworks, explaining exactly why our industry looks the way it does, and how all this growth—and more—is inevitable as the divide between man and machine blurs into oblivion.

The Gyrfalcon Eugene Potapov 2010-01-29 This book is the first monograph on one of the most beautiful and admired birds in the world. The Gyrfalcon is the world's largest and most powerful falcon - a truly awe-inspiring bird which inhabits the ferociously inhospitable Arctic taiga, from Greenland and Iceland right across Siberia and northern Canada. Its plumage varies from a dark mottled grey to pure white - the white birds in particular are coveted by birders and falconers. Like other titles in the series, it covers all aspects of the species' biology, taxonomy, distribution, status and historical associations with mankind. The result is an exhaustively researched and enthrallingly readable biography of a spectacular bird, illustrated throughout with photographs and line drawings. "a good read, presenting much previously inaccessible information, as well as some interesting insights into some of the most remote and inhospitable places on earth." British Birds

Trino: The Definitive Guide Matt Fuller 2021-04-14 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

Semantic Software Design Eben Hewitt 2019-09-25 With this practical book, architects, CTOs, and CIOs will learn a set of patterns for the practice of architecture, including analysis, documentation, and communication.

Author Eben Hewitt shows you how to create holistic and thoughtful technology plans, communicate them clearly, lead people toward the vision, and become a great architect or Chief Architect. This book covers each key aspect of architecture comprehensively, including how to incorporate business architecture, information architecture, data architecture, application (software) architecture together to have the best chance for the system's success. Get a practical set of proven architecture practices focused on shipping great products using architecture Learn how architecture works effectively with development teams, management, and product management teams through the value chain Find updated special coverage on machine learning architecture Get usable templates to start incorporating into your teams immediately Incorporate business architecture, information architecture, data architecture, and application (software) architecture together

Data Pipelines with Apache Airflow Julian de Ruyter 2021-04-05 "An Airflow bible. Useful for all kinds of users, from novice to expert." - Rambabu Posa, Sai Aashika Consultancy Data Pipelines with Apache Airflow teaches you how to build and maintain effective data pipelines. A successful pipeline moves data efficiently, minimizing pauses and blockages between tasks, keeping every process along the way operational. Apache Airflow provides a single customizable environment for building and managing data pipelines, eliminating the need for a hodgepodge collection of tools, snowflake code, and homegrown processes. Using real-world scenarios and examples, Data Pipelines with Apache Airflow teaches you how to simplify and automate data pipelines, reduce operational overhead, and smoothly integrate all the technologies in your stack. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Data pipelines manage the flow of data from initial collection through consolidation, cleaning, analysis, visualization, and more. Apache Airflow provides a single platform you can use to design, implement, monitor, and maintain your pipelines. Its easy-to-use UI, plug-and-play options, and flexible Python scripting make Airflow perfect for any data management task. About the book Data Pipelines with Apache Airflow teaches you how to build and maintain effective data pipelines. You'll explore the most common usage patterns, including aggregating multiple data sources, connecting to and from data lakes, and cloud deployment. Part reference and part tutorial, this practical guide covers every aspect of the directed acyclic graphs (DAGs) that power Airflow, and how to customize them for your pipeline's needs. What's inside Build, test, and deploy Airflow pipelines as DAGs Automate moving and transforming data Analyze historical datasets using backfilling Develop custom components Set up Airflow in production environments About the reader For DevOps, data engineers, machine learning engineers, and sysadmins with intermediate Python skills. About the author Bas Harens and Julian de Ruyter are data engineers with extensive experience using Airflow to develop pipelines for major companies. Bas is also an Airflow committer. Table of Contents PART 1 - GETTING STARTED 1 Meet Apache Airflow 2 Anatomy of an Airflow DAG 3 Scheduling in Airflow 4 Templating tasks using the Airflow context 5 Defining dependencies between tasks PART 2 - BEYOND THE BASICS 6 Triggering workflows 7 Communicating with external systems 8 Building custom components 9 Testing 10 Running tasks in containers PART 3 - AIRFLOW IN PRACTICE 11 Best practices 12 Operating Airflow in production 13 Securing Airflow 14 Project: Finding the fastest way to get around NYC PART 4 - IN THE CLOUDS 15 Airflow in the clouds 16 Airflow on AWS 17 Airflow on Azure 18 Airflow in GCP

Difficult Decisions in Surgical Ethics Vassyl A. Lonchyna 2022-01-11 This book provides a detailed guide to the ethical considerations involved when making decisions in surgery. Chapters feature a uniform format, which feature a case that represents a real-life problem, discussion of the medical indications of that issue, the latest available medical solutions, and related ethical considerations. In some cases, more in-depth debate is provided on why a particular decision should or should not be made based-upon ethical principles. Information boxes containing key statements and relevant data in clear easy-to-digest tables facilitates the reader in being able to assimilate the most important points covered in each chapter. Difficult Decisions in Surgical Ethics: An Evidence-Based Approach is a thorough review of ethical considerations in a range of surgical scenarios encompassing both adult and pediatric topics, training surgical residents, ethical care during a pandemic, critical care, palliative care, sensitivity to religious and ethnic mores, clinical research, and innovation. It is intended to be a vital resource for practicing and trainee surgeons seeking a comprehensive up-to-date resource on ethical topics in surgical practice. The work is part of the Difficult Decisions in Surgery series covering a range of surgical specialties.

The Crystal Bible Judy Hall 2016-04-01 A comprehensive and beautifully illustrated guide to crystals. Find a known crystal instantly or identify an unknown crystal in this easy-to-follow directory that includes photographic identification, detailed descriptions, and information on the individual properties of each crystal--including the spiritual, mental and psychological, emotional and physical effects, plus its use in healing. It's an indispensable reference for crystal lovers everywhere.

Cassandra: the Definitive Guide Louie Bryant 2015-08-19

Cassandra Kieran L. Lee 2015-08-12 Thought-provoking and accessible in approach, this updated and expanded second edition of the Cassandra: The Definitive Guide provides a user-friendly introduction to the subject, Taking a clear structural framework, it guides the reader through the subject's core elements. A flowing writing style combines with the use of illustrations and diagrams throughout the text to ensure the reader understands even the most complex of concepts. This succinct and enlightening overview is a required reading for advanced graduate-level students. We hope you find this book useful in shaping your future career. Feel free to send us your enquiries related to our publications to info@risepress.pw Rise Press

Technology Strategy Patterns Eben Hewitt 2018-10-15 Technologists who want their ideas heard, understood, and funded are often told to speak the language of business—without really knowing what that is. This book's toolkit provides architects, product managers, technology managers, and executives with a shared language—in the form of repeatable, practical patterns and templates—to produce great technology strategies. Author Eben Hewitt developed 39 patterns over the course of a decade in his work as CTO, CIO, and chief architect for several global tech companies. With these proven tools, you can define, create, elaborate, refine, and communicate your architecture goals, plans, and approach in a way that executives can readily understand, approve, and execute. This book covers: Architecture and strategy: Adopt a strategic architectural mindset to make a meaningful material impact Creating your strategy: Define the components of your technology strategy using proven patterns Communicating the strategy: Convey your technology strategy in a compelling way to a variety of audiences Bringing it all together: Employ patterns individually or in clusters for specific problems; use the complete framework for a comprehensive strategy

Cassandra Admin DBA Guide Prasad Bagewadi 2019-08-15 What value this book provides? This book absolutely provides tremendous value in terms its usefulness. This book takes away the pain associated with learning and mastering Cassandra. All complexity has been digested by the author and simplified for the reader with very useful and practical information that can be absorbed just by glancing through the pages. Years of author's experience and practical knowledge in Cassandra has been gifted to the reader in this book with great diligence and generosity. If you are planning to undergo expensive Cassandra training, think again, spending few hours with this book will change your mind, this book has been written with great care to reduce the learning curve. The aim of this book is multi fold, just to mention a few: Provide confidence to the reader in Cassandra concepts and architecture Provide a flexible, practical framework and context customizable for various situations Provide practical guidance to manage Cassandra platforms of various hues, sizes, shapes Provide real world examples to reduce guess work Provide executable query statements and command line statements at every step Provide practical outcomes to help the reader to gain instance understanding of what to expect Provide insights into making Cassandra environment robust and scalable Provide tricks and tips to implement and maintain seamlessly Provide security and vulnerability mitigation tips and steps Provide best practices to follow for optimal Cassandra use There is no doubt - this book makes the reader very productive Cassandra professional in very short span of time. This book essentially bridges the training gap as the industry is moving fast to take full advantage of what Cassandra can offer to fulfill emerging business needs. This book can be very helpful to Data administrators, Developers, Data modelers/Application Architects, Project Managers and Data Custodians. This book has range of topics, many are listed here: Cassandra concepts and architecture Cassandra Installation and Configuration Cassandra system architecture depicting gossip protocol, replication, consistency, tombstones, hinted handoff, compactions, repairs, memtables, commit log, read and write path functions Cassandra oriented data modelling Cassandra QL (CQL) tutorial Handling of Primary and Partition keys in Cassandra covering No joins, Static columns and TTL aspects Configuring authentication, authorization to access Cassandra in addition, steps to set up node-node and client-node SSL Configuring nodes addition, removal, decommission in single token and vnode setup modes in Cassandra Instructions to add new data center and delink the existing data center from a multi-dc cluster arrangement Cassandra backup and recovery functions with real examples of

restoring tables after truncation events
Cassandra utilities such as `sstabledump`, `sstablemetadata`, `sstablesplit`, `cqlsh` and `cassandra-stress`
Troubleshooting methods such as Node down, Read latency and Recover truncated table
Upgrading Cassandra to higher versions
Additional Cassandra architecture II methods such as Read and Write path, Compactions and Repairs
Cassandra High Availability Robbie Strickland 2014-12-29 If you are a developer or DevOps engineer who understands the basics of Cassandra and are ready to take your knowledge to the next level, then this book is for you. An understanding of the essentials of Cassandra is needed.

Mastering Apache Cassandra 3.x Aaron Ploetz 2018-10-31 Build, manage, and configure high-performing, reliable NoSQL database for your applications with Cassandra Key Features
Write programs more efficiently using Cassandra's features with the help of examples
Configure Cassandra and fine-tune its parameters depending on your needs
Integrate Cassandra database with Apache Spark and build strong data analytics pipeline
Book Description With ever-increasing rates of data creation, the demand for storing data fast and reliably becomes a need. Apache Cassandra is the perfect choice for building fault-tolerant and scalable databases. Mastering Apache Cassandra 3.x teaches you how to build and architect your clusters, configure and work with your nodes, and program in a high-throughput environment, helping you understand the power of Cassandra as per the new features. Once you've covered a brief recap of the basics, you'll move on to deploying and monitoring a production setup and optimizing and integrating it with other software. You'll work with the advanced features of CQL and the new storage engine in order to understand how they function on the server-side. You'll explore the integration and interaction of Cassandra components, followed by discovering features such as token allocation algorithm, CQL3, vnodes, lightweight transactions, and data modelling in detail. Last but not least you will get to grips with Apache Spark. By the end of this book, you'll be able to analyse big data, and build and manage high-performance databases for your application. What you will learn
Write programs more efficiently using Cassandra's features more efficiently
Exploit the given infrastructure, improve performance, and tweak the Java Virtual Machine (JVM)
Use CQL3 in your application in order to simplify working with Cassandra
Configure Cassandra and fine-tune its parameters depending on your needs
Set up a cluster and learn how to scale it
Monitor a Cassandra cluster in different ways
Use Apache Spark and other big data processing tools
Who this book is for Mastering Apache Cassandra 3.x is for you if you are a big data administrator, database administrator, architect, or developer who wants to build a high-performing, scalable, and fault-tolerant database. Prior knowledge of core concepts of databases is required.

Cassandra High Performance Cookbook Edward Capriolo 2011-07-15 Over 150 recipes to design and optimize large scale Apache Cassandra deployments.

Spark: The Definitive Guide Bill Chambers 2018-02-08 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

Cassandra: The Definitive Guide, (Revised) Third Edition Jeff Carpenter 2022-01-24 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

Cassandra: the Definitive Guide, (Revised) Third Edition Jeff Carpenter 2022-02 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

Cassandra Jeff Carpenter 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.

Leadership Through A Screen Joseph Brady 2019-04-04 The book defines and helps provide key solutions for some of the greatest leadership challenges facing global managers today. Leadership Through the Screen is a business leadership guidebook that tells a story. Written in an easy-to-read manner, each chapter highlights a single issue through the eyes of a fictional VP of marketing. The authors have done the research and included it in these pages so that business leaders do not have to. This book is meant to serve as a map to help modern managers weave their way through many of the fundamental challenges of leading people in a global and virtual realm. It provides the tools, knowledge, and potential solutions these leaders can use to forge successful and productive virtual teams.

The Fortune-Teller's Bible Jane Struthers 2007 A practical handbook designed to teach readers how to develop their divination skills looks at a wide variety of fortune-telling devices and techniques, including tarot, the I Ching, and scrying, and features skill-building exercises and instruction to teach newcomers how to use both Western and Chinese astrology, numerology, palmistry, runes, and more to forecast the future. Original.