Cassandra the definitive guide 2nd edition

Imagine a world where scalability isn't a problem, and you can focus on other aspects of your project. With this guide, you'll learn how to harness the power of Cassandra database management system to handle large amounts of data while maintaining high availability across multiple data centers. The second edition has been updated for Cassandra 3.0, providing you with technical details and practical examples to deploy this database in a production environment. Authors Jeff Carpenter and Eben Hewitt will guide you through the advantages of Cassandra's non-relational design, especially data modeling. This guide is perfect for developers, DBAs, or application architects who want to solve scalability issues or future-proof their applications. You'll learn how to use Cassandra's distributed and decentralized structure, create a working data model, develop sample applications, and deploy Cassandra on different platforms. With this comprehensive guide, you'll gain full access to Cassandra: The Definitive Guide, 2nd Edition, and get a free 10-day trial of O'Reilly. You can also take advantage of live events, courses curated by job role, and more. Cassandra for the Curious: A Technical Guide for Architects, Admins, and Developers For those seeking to understand the benefits of Cassandra, a NoSQL database system that deviates from traditional relational or object-oriented thinking. The book aims to provide a comprehensive guide for professionals and students alike, covering topics such as data modeling, implementation, and usage. Target audiences include: * Application architects and data architects exploring high-performance, decentralized, elastic data stores * Database administrators and developers transitioning to non-relational database systems * Managers making informed decisions on technology strategy involving Cassandra and related columnar databases * Students designing projects related to NoSQL databases This book assumes a solid understanding of Java programming and the Web development ecosystem. While Java knowledge is not mandatory, it enhances comprehension of exceptions, source code, and popular clients. Key features include: * Modular chapters designed as standalone guides for different skill levels and learning styles * Emphasis on ease of use and flexibility through various programming languages (C#, Python, node.js, PHP, Ruby) * Support for occasional self-setup and exploration of tools like Apache Ant and Git By following the book's structure, readers can navigate topics in a logical and approachable manner, regardless of their level of experience with Cassandra. Getting started with Cassandra involves running it and testing its basic features. Chapters 4-14 delve into various aspects of Cassandra. Given text here Examples from O'Reilly books require permission for reuse. Quoting code snippets or answering questions doesn't need permission, but incorporating significant amounts of example code does. We appreciate attribution, including title, author, publisher, and ISBN, like "Cassandra: The Definitive Guide, Second Edition, by Jeff Carpenter. Copyright 2016 Jeff Carpenter, 978-1-491-93366-4." If you're unsure about fair use or permission, contact permissions@oreilly.com. Safari is a membership-based training platform with access to thousands of books and resources from over 250 publishers. Members can visit oreilly.com/safari for more information. We thank our technical reviewers, editors, and contributors who helped bring this book to life. Imagine a world where scalability isn't an issue. This hands-on guide shows how Cassandra handles massive data while remaining highly available across multiple data centers. The expanded second edition provides the technical details and practical examples you need to put Cassandra to work in production. Cassandra Data Modeling Guide Overview The guide provides a comprehensive resource for developers, DBAs, and application architects looking to solve database scaling issues or future-proof their applications using Cassandra. The book helps users harness Cassandra's speed and flexibility by providing hands-on experience with CQL, data modeling, cluster topology, and integration with other tools like Spark and Presto.