

## Nosql And Sql Data Modeling Bringing Together Data Semantics And Software

Design great databases—from logical data modeling through physical schema definition. You will learn a framework that finally cracks the problem of merging data and process models into a meaningful and unified design that accounts for how data is actually used in production systems. Key to the framework is a method for taking the logical data model that is a static look at the definition of the data, and merging that static look with the process models describing how the data will be used in actual practice once a given system is implemented. The approach solves the disconnect between the static definition of data in the logical data model and the dynamic flow of the data in the logical process models. The design framework in this book can be used to create operational databases for transaction processing systems, or for data warehouses in support of decision support systems. The information manager can be a flat file, Oracle Database, IMS, NoSQL, Cassandra, Hadoop, or any other DBMS. Usage-Driven Database Design emphasizes practical aspects of design, and speaks to what works, what doesn't work, and what to avoid at all costs. Included in the book are lessons learned by the author over his 30+ years in the corporate trenches. Everything in the book is grounded on good theory, yet demonstrates a professional and pragmatic approach to design that can come only from decades of experience. Presents an end-to-end framework from logical data modeling through physical schema definition. Includes lessons learned, techniques, and tricks that can turn a database disaster into a success. Applies to all types of database management systems, including NoSQL such as Cassandra and Hadoop, and mainstream SQL databases such as Oracle and SQL Server What You'll Learn Create logical data models that accurately reflect the real world of the user Create usage scenarios reflecting how applications will use a new database Merge static data models with dynamic process models to create resilient yet flexible database designs Support application requirements by creating responsive database schemas in any database architecture Cope with big data and unstructured data for transaction processing and decision support systems Recognize when relational approaches won't work, and when to turn toward NoSQL solutions such as Cassandra or Hadoop Who This Book Is For System developers, including business analysts, database designers, database administrators, and application designers and developers who must design or interact with database systems

How do we design for data when traditional design techniques cannot extend to new database technologies? In this era of big data and the Internet of Things, it is essential that we have the tools we need to understand the data coming to us faster than ever before, and to design databases and data processing systems that can adapt easily to ever-changing data schemas and ever-changing business requirements. There must be no intellectual disconnect between data and the software that manages it. It must be possible to extract meaning and knowledge from data to drive artificial intelligence applications. Novel NoSQL data organization techniques must be used side-by-side with traditional SQL databases. Are existing data modeling techniques ready for all of this? The Concept and Object Modeling Notation (COMN) is able to cover the full spectrum of analysis and design. A single COMN model can represent the objects and concepts in the problem space, logical data design, and concrete NoSQL and SQL document, key-value, columnar, and relational database implementations. COMN models enable an unprecedented level of traceability of requirements to implementation. COMN models can also represent the static structure of software and the predicates that represent the patterns of meaning in databases. This book will teach you: the simple and familiar graphical notation of COMN with its three basic shapes and four line styles how to think about objects, concepts, types, and classes in the real world, using the ordinary meanings of English words that aren't tangled with confused techno-speak how to express logical data designs that are freer from implementation considerations than is possible in any other notation how to understand key-value, document, columnar, and table-oriented database designs in logical and physical terms how to use COMN to specify physical database implementations in any NoSQL or SQL database with the precision necessary for model-driven development

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems. Learn how different organizations are using graph databases to outperform their competitors. With this book's data modeling, query, and code examples, you'll quickly be able to implement your own solution. Model data with the Cypher query language and property graph model Learn best practices and common pitfalls when modeling with graphs Plan and implement a graph database solution in test-driven fashion Explore real-world examples to learn how and why organizations use a graph database Understand common patterns and components of graph database architecture Use analytical techniques and algorithms to mine graph database information

Advanced data management has always been at the core of efficient database and information systems. Recent trends like big data and cloud computing have aggravated the need for sophisticated and flexible data storage and processing solutions. This book provides a comprehensive coverage of the principles of data management developed in the last decades with a focus on data structures and query languages. It treats a wealth of different data models and surveys the foundations of structuring, processing, storing and querying data according these models. Starting off with the topic of database design, it further discusses weaknesses of the relational data model, and then proceeds to convey the basics of graph data, tree-structured XML data, key-value pairs and nested, semi-structured JSON data, columnar and record-oriented data as well as object-oriented data. The final chapters round the book off with an analysis of fragmentation, replication and consistency strategies for data management in distributed databases as well as recommendations for handling polyglot persistence in multi-model databases and multi-database architectures. While primarily geared towards students of Master-level courses in Computer Science and related areas, this book may also be of benefit to practitioners looking for a reference book on data modeling and query processing. It provides both theoretical depth and a concise treatment of open source technologies currently on the market.

This book constitutes the refereed proceedings of the 32nd International Conference on Conceptual Modeling, ER 2014, held in Atlanta, GA, USA. The 23 full and 15 short papers presented were carefully reviewed and selected from 80 submissions. Topics of interest presented and discussed in the conference span the entire spectrum of conceptual modeling including research and practice in areas such as: data on the web, unstructured data, uncertain and incomplete data, big data, graphs and networks, privacy and safety, database design, new modeling languages and applications, software concepts and strategies, patterns and narratives, data management for enterprise architecture, city and urban applications.

Data Modeling Essentials, Third Edition, covers the basics of data modeling while focusing on developing a facility in techniques, rather than a simple familiarization with "the rules". In order to enable students to apply the basics of data modeling to real models, the book addresses the realities of developing systems in real-world situations by assessing the merits of a variety of possible solutions as well as using language and diagramming methods that represent industry practice. This revised edition has been given significantly expanded coverage and reorganized for greater reader comprehension even as it retains its distinctive hallmarks of readability and usefulness. Beginning with the basics, the book provides a thorough grounding in theory before guiding the reader through the various stages of applied data modeling and database design. Later chapters address advanced subjects, including business rules, data warehousing, enterprise-wide modeling and data management. It includes an entirely new section discussing the development of logical and physical modeling, along with new material describing a powerful technique for model verification. It also provides an excellent resource for additional lectures and exercises. This text is the ideal reference for data modelers, data architects, database designers, DBAs, and systems analysts, as well as undergraduate and graduate-level students looking for a real-world perspective. Thorough coverage of the fundamentals and relevant theory. Recognition and support for the creative side of the process. Expanded coverage of applied data modeling includes new chapters on logical and physical database design. New material describing a powerful technique for model verification. Unique coverage of the practical and human aspects of modeling, such as working with business specialists, managing change, and resolving conflict.

NoSQL Starter is a great resource for someone starting with NoSQL and an indispensable guide for technology decision makers. It is assumed that you have a background in RDBMS modeling and SQL and have had exposure to at least one of the programming languages – Java or JavaScript. Friendly, practical tutorial with lots of hints and tips from several experienced Solr users and developers.

Big Data Analytics will assist managers in providing an overview of the drivers for introducing big data technology into the organization and for understanding the types of business problems best suited to big data analytics solutions, understanding the value drivers and benefits, strategic planning, developing a pilot, and eventually planning to integrate back into production within the enterprise. Guides the reader in assessing the opportunities and value proposition Overview of big data hardware and software architectures Presents a variety of technologies and how they fit into the big data ecosystem

NoSQL was developed to overcome the limitations of relational databases in the largest Web applications at companies such as Google, Yahoo and Facebook. As it is applied more widely, developers are finding that it can simplify scalability while requiring far less coding and management overhead. However, NoSQL requires fundamentally different approaches to database design and modeling, and many conventional relational techniques lead to suboptimal results. *NoSQL for Mere Mortals* is an easy, practical guide to succeeding with NoSQL in your environment. Following the classic, best-selling format pioneered in *SQL Queries for Mere Mortals*, enterprise database expert Dan Sullivan guides you step-by-step through choosing technologies, designing high-performance databases, and planning for long-term maintenance. *Sullivan* introduces each type of NoSQL database, shows how to install and manage them, and demonstrates how to leverage their features while avoiding common mistakes that lead to poor performance and unmet requirements. He uses four popular NoSQL databases as reference models: MongoDB, a document database; Cassandra, a column family data store; Redis, a key-value database; and Neo4j, a graph database. You'll find explanations of each database's structure and capabilities, practical guidelines for choosing amongst them, and expert guidance on designing databases with them. *Packed with examples, NoSQL for Mere Mortals is today's best way to master NoSQL -- whether you're a DBA, developer, user, or student.*

Master a graph data modeling technique superior to traditional data modeling for both relational and NoSQL databases (graph, document, key-value, and column), leveraging cognitive psychology to improve big data designs. From Karen Lopez's Foreword: In this book, Thomas Frisendal raises important questions about the continued usefulness of traditional data modeling notations and approaches: Are Entity Relationship Diagrams (ERDs) relevant to analytical data requirements? Are ERDs relevant in the new world of Big Data? Are ERDs still the best way to work with business users to understand their needs? Are Logical and Physical Data Models too closely coupled? Are we correct in using the same notations for communicating with business users and developers? Should we refine our existing notations and tools to meet these new needs, or should we start again from a blank page? What new notations and approaches will we need? How will we use those to build enterprise database systems? Frisendal takes us through the history of data modeling, enterprise data models and traditional modeling methods. He points out, quite contentiously, where he feels we have gone wrong and in a few places where we got it right. He then maps out the psychology of meaning and context, while identifying important issues about where data modeling may or may not fit in business modeling. The main subject of this work is a proposal for a new exploration-driven modeling approach and new modeling notations for business concept models, business solutions models, and physical data models with examples on how to leverage those for implementing into any target database or datastore. These new notations are based on a property graph approach to modeling data.

Joe Celko's Complete Guide to NoSQL provides a complete overview of non-relational technologies so that you can become more nimble to meet the needs of your organization. As data continues to explode and grow more complex, SQL is becoming less useful for querying data and extracting meaning. In this new world of bigger and faster data, you will need to leverage non-relational technologies to get the most out of the information you have. Learn where, when, and why the benefits of NoSQL outweigh those of SQL with Joe Celko's Complete Guide to NoSQL. This book covers three areas that make today's new data different from the data of the past: velocity, volume and variety. When information is changing faster than you can collect and query it, it simply cannot be treated the same as static data. Celko will help you understand velocity, to equip you with the tools to drink

from a fire hose. Old storage and access models do not work for big data. Celko will help you understand volume, as well as different ways to store and access data such as petabytes and exabytes. Not all data can fit into a relational model, including genetic data, semantic data, and data generated by social networks. Celko will help you understand variety, as well as the alternative storage, query, and management frameworks needed by certain kinds of data. Gain a complete understanding of the situations in which SQL has more drawbacks than benefits so that you can better determine when to utilize NoSQL technologies for maximum benefit Recognize the pros and cons of columnar, streaming, and graph databases Make the transition to NoSQL with the expert guidance of best-selling SQL expert Joe Celko

The Concept and Object Modeling Notation (COMN) is able to cover the full spectrum of analysis and design. A single COMN model can represent the objects and concepts in the problem space, logical data design, and concrete NoSQL and SQL document, key-value, columnar, and relational database implementations. COMN models enable an unprecedented level of traceability of requirements to implementation. COMN models can also represent the static structure of software and the predicates that represent the patterns of meaning in databases.

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

Introductory, theory-practice balanced text teaching the fundamentals of databases to advanced undergraduates or graduate students in information systems or computer science.

Choose the right Azure data service and correct model design for successful implementation of your data model with the help of this hands-on guide Key Features Design a cost-effective, performant, and scalable database in Azure Choose and implement the most suitable design for a database Discover how your database can scale with growing data volumes, concurrent users, and query complexity Book Description Data is at the heart of all applications and forms the foundation of modern data-driven businesses. With the multitude of data-related use cases and the availability of different data services, choosing the right service and implementing the right design becomes paramount to successful implementation. Data Modeling for Azure Data Services starts with an introduction to databases, entity analysis, and normalizing data. The book then shows you how to design a NoSQL database for optimal performance and scalability and covers how to provision and implement Azure SQL DB, Azure Cosmos DB, and Azure Synapse SQL Pool. As you progress through the chapters, you'll learn about data analytics, Azure Data Lake, and Azure SQL Data Warehouse and explore dimensional modeling, data vault modeling, along with designing and implementing a Data Lake using Azure Storage. You'll also learn how to implement ETL with Azure Data Factory. By the end of this book, you'll have a solid understanding of which Azure data services are the best fit for your model and how to implement the best design for your solution. What you will learn Model relational database using normalization, dimensional, or Data Vault modeling Provision and implement Azure SQL DB and Azure Synapse SQL Pools Discover how to model a Data Lake and implement it using Azure Storage Model a NoSQL database and provision and implement an Azure Cosmos DB Use Azure Data Factory to implement ETL/ELT processes Create a star schema model using dimensional modeling Who this book is for This book is for business intelligence developers and consultants who work on (modern) cloud data warehousing and design and implement databases. Beginner-level knowledge of cloud data management is expected.

"It's not easy to find such a generous book on big data and databases. Fortunately, this book is the one." Feng Yu. Computing Reviews. June 28, 2016. This is a book for enterprise architects, database administrators, and developers who need to understand the latest developments in database technologies. It is the book to help you choose the correct database technology at a time when concepts such as Big Data, NoSQL and NewSQL are making what used to be an easy choice into a complex decision with significant implications. The relational database (RDBMS) model completely dominated database technology for over 20 years. Today this "one size fits all" stability has been disrupted by a relatively recent explosion of new database technologies. These paradigm-busting technologies are powering the "Big Data" and "NoSQL" revolutions, as well as forcing fundamental changes in databases across the board. Deciding to use a relational database was once truly a no-brainer, and the various commercial relational databases competed on price, performance, reliability, and ease of use rather than on fundamental architectures. Today we are faced with choices between radically different database technologies. Choosing the right database today is a complex undertaking, with serious economic and technological consequences. Next Generation Databases demystifies today's new database technologies. The book describes what each technology was designed to solve. It shows how each technology can be used to solve real word application and business problems. Most importantly, this book highlights the architectural differences between technologies that are the critical factors to consider when choosing a database platform for new and upcoming projects. Introduces the new technologies that have revolutionized the database landscape Describes how each technology can be

used to solve specific application or business challenges Reviews the most popular new wave databases and how they use these new database technologies

Data is getting bigger and more complex by the day, and so are your choices in handling it. Explore some of the most cutting-edge databases available - from a traditional relational database to newer NoSQL approaches - and make informed decisions about challenging data storage problems. This is the only comprehensive guide to the world of NoSQL databases, with in-depth practical and conceptual introductions to seven different technologies: Redis, Neo4J, CouchDB, MongoDB, HBase, Postgres, and DynamoDB. This second edition includes a new chapter on DynamoDB and updated content for each chapter. While relational databases such as MySQL remain as relevant as ever, the alternative, NoSQL paradigm has opened up new horizons in performance and scalability and changed the way we approach data-centric problems. This book presents the essential concepts behind each database alongside hands-on examples that make each technology come alive. With each database, tackle a real-world problem that highlights the concepts and features that make it shine. Along the way, explore five database models - relational, key/value, columnar, document, and graph - from the perspective of challenges faced by real applications. Learn how MongoDB and CouchDB are strikingly different, make your applications faster with Redis and more connected with Neo4J, build a cluster of HBase servers using cloud services such as Amazon's Elastic MapReduce, and more. This new edition brings a brand new chapter on DynamoDB, updated code samples and exercises, and a more up-to-date account of each database's feature set. Whether you're a programmer building the next big thing, a data scientist seeking solutions to thorny problems, or a technology enthusiast venturing into new territory, you will find something to inspire you in this book. What You Need: You'll need a \*nix shell (Mac OS or Linux preferred, Windows users will need Cygwin), Java 6 (or greater), and Ruby 1.8.7 (or greater). Each chapter will list the downloads required for that database.

A beginner's guide to get you up and running with Cassandra, DynamoDB, HBase, InfluxDB, MongoDB, Neo4j, and Redis Key Features Covers the basics of 7 NoSQL databases and how they are used in the enterprises Quick introduction to MongoDB, DynamoDB, Redis, Cassandra, Neo4j, InfluxDB, and HBase Includes effective techniques for database querying and management Book Description This is the golden age of open source NoSQL databases. With enterprises having to work with large amounts of unstructured data and moving away from expensive monolithic architecture, the adoption of NoSQL databases is rapidly increasing. Being familiar with the popular NoSQL databases and knowing how to use them is a must for budding DBAs and developers. This book introduces you to the different types of NoSQL databases and gets you started with seven of the most popular NoSQL databases used by enterprises today. We start off with a brief overview of what NoSQL databases are, followed by an explanation of why and when to use them. The book then covers the seven most popular databases in each of these categories: MongoDB, Amazon DynamoDB, Redis, HBase, Cassandra, InfluxDB, and Neo4j. The book doesn't go into too much detail about each database but teaches you enough to get started with them. By the end of this book, you will have a thorough understanding of the different NoSQL databases and their functionalities, empowering you to select and use the right database according to your needs. What you will learn Understand how MongoDB provides high-performance, high-availability, and automatic scaling Interact with your Neo4j instances via database queries, Python scripts, and Java application code Get familiar with common querying and programming methods to interact with Redis Study the different types of problems Cassandra can solve Work with HBase components to support common operations such as creating tables and reading/writing data Discover data models and work with CRUD operations using DynamoDB Discover what makes InfluxDB a great choice for working with time-series data Who this book is for If you are a budding DBA or a developer who wants to get started with the fundamentals of NoSQL databases, this book is for you. Relational DBAs who want to get insights into the various offerings of popular NoSQL databases will also find this book to be very useful.

Big data modeling is very challenging to handle using traditional database modeling and management systems. This book will teach you how to model big data using the latest and more efficient tools such as ERWIN, ANACONDA (Python), and WEKA to model data.

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

Summary Making Sense of NoSQL clearly and concisely explains the concepts, features, benefits, potential, and limitations of NoSQL technologies. Using examples and use cases, illustrations, and plain, jargon-free writing, this guide shows how you can effectively assemble a NoSQL solution to replace or augment the traditional RDBMS you have now. About this Book If you want to understand and perhaps start using the new data storage and analysis technologies that go beyond the SQL database model, this book is for you. Written in plain language suitable for technical managers and developers, and using many examples, use cases, and illustrations, this book explains the concepts, features, benefits, potential, and limitations of NoSQL. Making Sense of NoSQL starts by comparing familiar database concepts to the new NoSQL patterns that augment or replace them. Then, you'll explore case studies on big data, search, reliability, and business agility that apply these new patterns to today's business problems. You'll see how NoSQL systems can leverage the resources of modern cloud computing and multiple-CPU data centers. The final chapters show you how to choose the right NoSQL technologies for your own needs. Managers and developers will welcome this lucid overview of the potential and capabilities of NoSQL technologies. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside NoSQL data architecture patterns NoSQL for big data Search, high availability, and security Choosing an architecture About the Authors Dan McCreary and Ann Kelly lead an independent training and consultancy firm focused on NoSQL solutions and are cofounders of the NoSQL Now! Conference. Table of Contents PART 1 INTRODUCTION NoSQL: It's about making intelligent choices NoSQL concepts PART 2 DATABASE PATTERNS Foundational data architecture patterns NoSQL data architecture patterns Native XML databases PART 3 NOSQL SOLUTIONS Using NoSQL to manage big data Finding information with NoSQL search Building high-availability solutions with NoSQL Increasing agility with NoSQL PART 4 ADVANCED TOPICS NoSQL and functional programming Security: protecting data in your NoSQL systems Selecting the right NoSQL solution

This book offers a comprehensive introduction to relational (SQL) and non-relational (NoSQL) databases. The authors thoroughly review the current state of database tools and techniques, and examine coming innovations. The book opens with a broad look at data management, including an overview of information systems and databases, and an explanation of contemporary database types: SQL and NoSQL databases, and their respective management systems The nature and uses of Big Data A high-level view of the organization of data management Data Modeling and Consistency Chapter-length treatment is afforded Data Modeling in both relational and graph databases, including enterprise-wide data architecture, and formulas for database design. Coverage of languages extends from an overview of operators, to SQL and and QBE (Query by Example), to integrity constraints and more. A full chapter probes the challenges of Ensuring Data Consistency, covering: Multi-User Operation Troubleshooting

Consistency in Massive Distributed Data Comparison of the ACID and BASE consistency models, and more System Architecture also gets from its own chapter, which explores Processing of Homogeneous and Heterogeneous Data; Storage and Access Structures; Multi-dimensional Data Structures and Parallel Processing with MapReduce, among other topics. Post-Relational and NoSQL Databases The chapter on post-relational databases discusses the limits of SQL – and what lies beyond, including Multi-Dimensional Databases, Knowledge Bases and and Fuzzy Databases. A final chapter covers NoSQL Databases, along with Development of Non-Relational Technologies, Key-Value, Column-Family and Document Stores XML Databases and Graphic Databases, and more The book includes more than 100 tables, examples and illustrations, and each chapter offers a list of resources for further reading. SQL & NoSQL Databases conveys the strengths and weaknesses of relational and non-relational approaches, and shows how to undertake development for big data applications. The book benefits readers including students and practitioners working across the broad field of applied information technology. This textbook has been recommended and developed for university courses in Germany, Austria and Switzerland.

Get up to speed on the nuances of NoSQL databases and what they mean for your organization This easy to read guide to NoSQL databases provides the type of no-nonsense overview and analysis that you need to learn, including what NoSQL is and which database is right for you. Featuring specific evaluation criteria for NoSQL databases, along with a look into the pros and cons of the most popular options, NoSQL For Dummies provides the fastest and easiest way to dive into the details of this incredible technology. You'll gain an understanding of how to use NoSQL databases for mission-critical enterprise architectures and projects, and real-world examples reinforce the primary points to create an action-oriented resource for IT pros. If you're planning a big data project or platform, you probably already know you need to select a NoSQL database to complete your architecture. But with options flooding the market and updates and add-ons coming at a rapid pace, determining what you require now, and in the future, can be a tall task. This is where NoSQL For Dummies comes in! Learn the basic tenets of NoSQL databases and why they have come to the forefront as data has outpaced the capabilities of relational databases Discover major players among NoSQL databases, including Cassandra, MongoDB, MarkLogic, Neo4J, and others Get an in-depth look at the benefits and disadvantages of the wide variety of NoSQL database options Explore the needs of your organization as they relate to the capabilities of specific NoSQL databases Big data and Hadoop get all the attention, but when it comes down to it, NoSQL databases are the engines that power many big data analytics initiatives. With NoSQL For Dummies, you'll go beyond relational databases to ramp up your enterprise's data architecture in no time. This book is intended for database professionals, software developers, and architects who have some previous experience with MongoDB and now want to shift their focus to the concepts of data modeling. If you wish to develop better schema designs for MongoDB-based applications, this book is ideal for you.

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

? An important step in database implementation is the data modeling, because it facilitates the understanding of the project through key features that can prevent programming and operation errors. ? In database technologies, some of the new issues increasingly debated arenon-conventional applications, including NoSQL (Not only SQL) databases, whichwere initially created in response to the needs for better scalability, lowerlatency and higher flexibility in an era of bigdata and cloud computing. Thesenon-functional aspects are the main reason for using NoSQL database. ? Data modeling has an important role to play in NoSQL environments. The datamodeling process involves the creation of a diagram that represents the meaning of the data and the relationship between the data elements. Thus, understanding is a fundamental aspect of data modeling and a pattern for this kind of representation has few contributions for NoSQL databases. ? This book explains a NoSQL data modeling standard, introducing modeling techniques that can be used on document-oriented databases. We have considered Cassandra and Riak NoSQL databases because of the heterogeneous characteristics of each NoSQL database classification so that to fill the knowledge gap by studying the available non-relational databases in order to develop a systematic approach for solving problems of data persistence using these technologies. Ajit & Sultan.....

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.

What value does semantic data modeling offer? As an information architect or data science professional, let's say you have an abundance of the right data and the technology to extract business gold—but you still fail. The reason? Bad data semantics. In this practical and comprehensive field guide, author Panos Alexopoulos takes you on an eye-opening journey through semantic data modeling as applied in the real world. You'll learn how to master this craft to increase the usability and value of your data and applications. You'll also explore the pitfalls to avoid and dilemmas to overcome for building high-quality and valuable semantic representations of data. Understand the fundamental concepts, phenomena, and processes related to semantic data modeling Examine the quirks and challenges of semantic data modeling and learn how to effectively leverage the available frameworks and tools Avoid mistakes and bad practices that can undermine your efforts to create good data models Learn about model development dilemmas, including representation, expressiveness and content, development, and governance Organize and execute semantic data initiatives in your organization, tackling technical, strategic, and organizational challenges

?This book undertakes to marry the concepts of "Concept Mapping" with a "Design Thinking" approach in the context of business analysis. While in the past a lot of attention has been paid to the business process side, this book now focusses information quality and valuation, master data and hierarchy management, business rules automation and business semantics as examples for business innovation opportunities. The book shows how to take "Business Concept Maps" further as information models for new IT paradigms. In a way this books redefines and extends business analysis towards solutions that can be described as business synthesis or business development. Business modellers, analysts and controllers, as well as enterprise information architects, will benefit from the intuitive modelling and designing approach presented in this book. The pragmatic and agile methods presented can be directly applied to improve the way organizations manage their business concepts and their relationships. "This book is a great contribution to the information management community. It combines a theoretical foundation with practical methods for dealing with important problems. This is rare and very useful.

Conceptual models that communicate business reality effectively require some degree of creative imagination. As such, they combine the results of business analysis with communication design, as is extensively covered in this book." Dr. Malcolm Chisholm, President at AskGet.com Inc. "Truly understanding business requirements has always been a major stumbling block in business intelligence (BI) projects. In this book, Thomas Frisendal introduces a powerful technique—business concept mapping—that creates a virtual mind-meld between business users and business analysts. Frisendal does a wonderful explaining and demonstrating how this tool can improve the outcome of BI and other development projects ." Wayne Eckerson, executive director, BI Leadership Forum

This book constitutes the proceedings of the 11th European Workshop on Performance Engineering, EPEW 2014, held in Florence, Italy, in September 2014. The 18 full papers presented in this volume were carefully reviewed and selected from 30 submissions. The papers are organized in topical sections named: cloud performance modelling; queueing and fluid models; performance of computation and programming; fitting; urban traffic modelling; decision making; and Markovian models, above and beyond.

Data Modeling Made Simple will provide the business or IT professional with a practical working knowledge of data modeling concepts and best practices. This book is written in a conversational style that encourages you to read it from start to finish and master these ten objectives: Know when a data model is needed and which type of data model is most effective for each situation Read a data model of any size and complexity with the same confidence as reading a book Build a fully normalized relational data model, as well as an easily navigatable dimensional model Apply techniques to turn a logical data model into an efficient physical design Leverage several templates to make requirements gathering more efficient and accurate Explain all ten categories of the Data Model Scorecard Learn strategies to improve your working relationships with others Appreciate the impact unstructured data has, and will have, on our data modeling deliverables Learn basic UML concepts Put data modeling in context with XML, metadata, and agile development Book Review by Johnny Gay In this book review, I address each section in the book and provide what I found most valuable as a data modeler. I compare, as I go, how the book's structure eases the new data modeler into the subject much like an instructor might ease a beginning swimmer into the pool. This book begins like a Dan Brown novel. It even starts out with the protagonist, our favorite data modeler, lost on a dark road somewhere in France. In this case, what saves him isn't a cipher, but of all things, something that's very much like a data model in the form of a map! The author deems they are both way-finding tools. The chapters in the book are divided into 5 sections. The chapters in each section end with an exercise and a list of the key points covered to reinforce what you've learned. I find myself comparing the teaching structure of the book to the way most of us learn to swim.

The topic of NoSQL databases has recently emerged, to face the Big Data challenge, namely the ever increasing volume of data to be handled. It is now recognized that relational databases are not appropriate in this context, implying that new database models and techniques are needed. This book presents recent research works, covering the following basic aspects: semantic data management, graph databases, and big data management in cloud environments. The chapters in this book report on research about the evolution of basic concepts such as data models, query languages, and new challenges regarding implementation issues.

The need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational "NoSQL" databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage's pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Polyglot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access.

This presentation was recorded live at Data Modeling Zone (DMZ) 2016. The venerable entity-relationship (E-R) modeling technique is very mature for SQL databases, but doesn't translate well to many of the NOSQL databases. This interactive workshop will introduce the Concept and Object Modeling Notation (COMN, pronounced "common"), which can represent the new structures that NOSQL DBMSs support, while still supporting traditional SQL DBMSs. It also has more expressive power for modeling the real-world entities that databases are about. Topics

to be covered include: modeling notation for containment and multi-valued attributes how to model schema-less databases how to model the problem space and the solution space, and the mapping between them.

Fully revised and updated, Relational Database Design, Second Edition is the most lucid and effective introduction to relational database design available. Here, you'll find the conceptual and practical information you need to develop a design that ensures data accuracy and user satisfaction while optimizing performance, regardless of your experience level or choice of DBMS. Supporting the book's step-by-step instruction are three case studies illustrating the planning, analysis, and design steps involved in arriving at a sound design. These real-world examples include object-relational design techniques, which are addressed in greater detail in a new chapter devoted entirely to this timely subject. \* Concepts you need to master to put the book's practical instruction to work. \* Methods for tailoring your design to the environment in which the database will run and the uses to which it will be put. \* Design approaches that ensure data accuracy and consistency. \* Examples of how design can inhibit or boost database application performance. \* Object-relational design techniques, benefits, and examples. \* Instructions on how to choose and use a normalization technique. \* Guidelines for understanding and applying Codd's rules. \* Tools to implement a relational design using SQL. \* Techniques for using CASE tools for database design.

Master a graph data modeling technique superior to traditional data modeling for both relational and NoSQL databases (graph, document, key-value, and column), leveraging cognitive psychology to improve big data designs. From Karen Lopez's Foreword: In this book, Thomas Frisendal raises important questions about the continued usefulness of traditional data modeling notations and approaches:  $\phi\phi$ Are Entity Relationship Diagrams (ERDs) relevant to analytical data requirements?  $\phi\phi$ Are ERDs relevant in the new world of Big Data?  $\phi\phi$ Are ERDs still the best way to work with business users to understand their needs?  $\phi\phi$ Are Logical and Physical Data Models too closely coupled?  $\phi\phi$ Are we correct in using the same notations for communicating with business users and developers?  $\phi\phi$ Should we refine our existing notations and tools to meet these new needs, or should we start again from a blank page?  $\phi\phi$ What new notations and approaches will we need?  $\phi\phi$ How will we use those to build enterprise database systems? Frisendal takes us through the history of data modeling, enterprise data models and traditional modeling methods. He points out, quite contentiously, where he feels we have gone wrong and in a few places where we got it right. He then maps out the psychology of meaning and context, while identifying important issues about where data modeling may or may not fit in business modeling. The main subject of this work is a proposal for a new exploration-driven modeling approach and new modeling notations for business concept models, business solutions models, and physical data models with examples on how to leverage those for implementing into any target database or datastore. These new notations are based on a property graph approach to modeling data.

If you are interested in Cassandra and want to develop real-world analysis applications, then this book is perfect for you. It would be helpful to have prior knowledge of NoSQL database. Graph data closes the gap between the way humans and computers view the world. While computers rely on static rows and columns of data, people navigate and reason about life through relationships. This practical guide demonstrates how graph data brings these two approaches together. By working with concepts from graph theory, database schema, distributed systems, and data analysis, you'll arrive at a unique intersection known as graph thinking. Authors Denise Koessler Gosnell and Matthias Broecheler show data engineers, data scientists, and data analysts how to solve complex problems with graph databases. You'll explore templates for building with graph technology, along with examples that demonstrate how teams think about graph data within an application. Build an example application architecture with relational and graph technologies Use graph technology to build a Customer 360 application, the most popular graph data pattern today Dive into hierarchical data and troubleshoot a new paradigm that comes from working with graph data Find paths in graph data and learn why your trust in different paths motivates and informs your preferences Use collaborative filtering to design a Netflix-inspired recommendation system

"The Concept and Object Modeling Notation (COMN) supports contextual, logical, and physical data modeling for NoSQL and SQL DBMSs. In this advanced class, Ted Hills will walk through two physical models for a hypothetical coffee shop business, exploring a design for a document database for order capture and a traditional SQL database for data warehousing (a Star Schema). A COMN logical model ties it all together, ensuring that the two very different physical database designs are representing the same logical data. Along the way, Ted will explain the COMN notation."--Resource description page.

[Copyright: 276aa80a8f59d5e5840a7ce4cf5ba1ce](https://www.coursera.org/learn/comn)