

Database Modeling And Design The Fundamental Principles The Morgan Kaufmann Series In Data Management Systems

This book is for all data modelers, data architects, and database designers?be they novices who want to learn what's involved in data modeling, or experienced modelers who want to brush up their skills. A novice will not only gain an overview of data modeling, they will also learn how to follow the data modeling process, including the activities required for each step. The experienced practitioner will discover (or rediscover) techniques to ensure that data models accurately reflect business requirements. This book describes rigorous yet easily implemented approaches to:

- modeling of business information requirements for review by business stakeholders before development of the logical data model
- normalizing data, based on simple questions rather than the formal definitions which many modelers find intimidating
- naming and defining concepts and attributes
- modeling of time-variant data
- documenting business rules governing both the real world and data
- data modeling in an Agile project
- managing data model change in any type of project
- transforming a business information model to a logical data model against which developers can code
- implementing the logical data model in a traditional relational DBMS, an SQL:2003-compliant DBMS, an object-relational DBMS, or in XML.

Part 1 describes business information models in-depth, including:

- the importance of modeling business information requirements before embarking on a logical data model
- business concepts (entity classes)
- attributes of business concepts
- attribute classes as an alternative to DBMS data types
- relationships between business concepts
- time-variant data
- generalization and specialization of business concepts
- naming and defining the components of the business information model
- business rules governing data, including a distinction between real-world rules and data rules.

Part 2 journeys from requirements to a working data resource, covering:

- sourcing data requirements
- developing the business information model
- communicating it to business stakeholders for review, both as diagrams and verbally
- managing data model change
- transforming the business information model into a logical data model of stored data for implementation in a relational or object-relational DBMS
- attribute value representation and data constraints (important but often overlooked)
- modeling data vault, dimensional and XML data.

Shows techniques for managing the complexity of database design using the ER model, a popular method for representing data requirements. Presents a complete set of semantic definitions and notations for ER models with computer screen illustrations of large, complex databases. Includes both logical and physical database design with an emphasis on the former. Annotation copyrighted by Book News, Inc., Portland, OR

DATA MODELING AND DATABASE DESIGN presents a conceptually complete coverage of indispensable topics that each MIS student should learn if that student takes only one database course. Database design and data modeling encompass the minimal set of topics addressing the core competency of knowledge students should acquire in the database area. The text, rich examples, and figures work together to cover material with a depth and precision that is not available in more introductory database books. Important Notice: Media content referenced within the product description or the product text may not be available in the ebook version.

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

Agile Data Warehouse Design is a step-by-step guide for capturing data warehousing/business intelligence (DW/BI) requirements and turning them into high performance dimensional models in the most direct way: by modelstorming (data modeling] brainstorming) with BI stakeholders. This book describes BEAM, an agile approach to dimensional modeling, for improving communication between data warehouse designers, BI stakeholders and the whole DW/BI development team. BEAM provides tools and techniques that will encourage DW/BI designers and developers to move away from their keyboards and entity relationship based tools and model interactively with their colleagues. The result is everyone thinks dimensionally from the outset! Developers understand how to efficiently implement dimensional modeling solutions. Business stakeholders feel ownership of the data warehouse they have created, and can already imagine how they will use it to answer their business questions. Within this book, you will learn: Agile dimensional modeling using Business Event Analysis & Modeling (BEAM) Modelstorming: data modeling that is quicker, more inclusive, more productive, and frankly more fun! Telling dimensional data stories using the 7Ws (who, what, when, where, how many, why and how) Modeling by example not abstraction; using data story themes, not crow's feet, to describe detail Storyboarding the data warehouse to discover conformed dimensions and plan iterative development Visual modeling: sketching timelines, charts and grids to model complex process measurement - simply Agile design documentation: enhancing star schemas with BEAM dimensional shorthand notation Solving difficult DW/BI performance and usability problems with proven dimensional design patterns LawrenceCorr is a data warehouse designer and educator. As Principal of DecisionOne Consulting, he helps clients to review and simplify their data warehouse designs, and advises vendors on visual data modeling techniques. He regularly teaches agile dimensional modeling courses worldwide and has taught dimensional DW/BI skills to thousands of students. Jim Stagnitto is a data warehouse and master data management architect specializing in the healthcare, financial services, and information service industries. He is the founder of the data warehousing and data mining consulting firm Llumino.

Essential Skills--Made Easy! Learn how to create data models that allow complex data to be analyzed, manipulated, extracted, and reported upon accurately. Data Modeling: A Beginner's Guide teaches you techniques for gathering business requirements and using them to produce conceptual, logical, and physical database designs. You'll get details on Unified Modeling Language (UML), normalization, incorporating business rules, handling temporal data, and analytical database design. The methods presented in this fast-paced tutorial are applicable to any database management system, regardless of vendor. Designed for Easy Learning Key Skills & Concepts--Chapter-opening lists of specific skills covered in the chapter Ask the expert--Q&A sections filled with bonus information and helpful tips Try

This--Hands-on exercises that show you how to apply your skills Notes--Extra information related to the topic being covered Self Tests--Chapter-ending quizzes to test your knowledge Andy Opper has taught database technology for the University of California Extension for more than 25 years. He is the author of Databases Demystified, SQL Demystified, and Databases: A Beginner's Guide, and the co-author of SQL: A Beginner's Guide, Third Edition, and SQL: The Complete Reference, Third Edition.

Information Modeling and Relational Databases provides an introduction to ORM (Object Role Modeling)-and much more. In fact, it's the only book to go beyond introductory coverage and provide all of the in-depth instruction you need to transform knowledge from domain experts into a sound database design. Inside, ORM authority Terry Halpin blends conceptual information with practical instruction that will let you begin using ORM effectively as soon as possible. Supported by examples, exercises, and useful background information, his step-by-step approach teaches you to develop a natural-language-based ORM model and then, where needed, abstract ER and UML models from it. This book will quickly make you proficient in the modeling technique that is proving vital to the development of accurate and efficient databases that best meet real business objectives. The most in-depth coverage of Object Role Modeling available anywhere-written by a pioneer in the development of ORM. Provides additional coverage of Entity Relationship (ER) modeling and the Unified Modeling Language-all from an ORM perspective. Intended for anyone with a stake in the accuracy and efficacy of databases: systems analysts, information modelers, database designers and administrators, instructors, managers, and programmers. Explains and illustrates required concepts from mathematics and set theory.

Database Modeling and Design, Fifth Edition, focuses on techniques for database design in relational database systems. This extensively revised fifth edition features clear explanations, lots of terrific examples and an illustrative case, and practical advice, with design rules that are applicable to any SQL-based system. The common examples are based on real-life experiences and have been thoroughly class-tested. This book is immediately useful to anyone tasked with the creation of data models for the integration of large-scale enterprise data. It is ideal for a stand-alone data management course focused on logical database design, or a supplement to an introductory text for introductory database management. In-depth detail and plenty of real-world, practical examples throughout Loaded with design rules and illustrative case studies that are applicable to any SQL, UML, or XML-based system Immediately useful to anyone tasked with the creation of data models for the integration of large-scale enterprise data. Database systems and database design technology have undergone significant evolution in recent years. The relational data model and relational database systems dominate business applications; in turn, they are extended by other technologies like data warehousing, OLAP, and data mining. How do you model and design your database application in consideration of new technology or new business needs? In the extensively revised fifth edition, you'll get clear explanations, lots of terrific examples and an illustrative case, and the really practical advice you have come to count on--with design rules that are applicable to any SQL-based system. But you'll also get plenty to help you grow from a new database designer to an experienced designer developing industrial-sized systems. In-depth detail and plenty of real-world, practical examples throughout Loaded with design rules and illustrative case studies that are applicable to any SQL, UML, or XML-based system Immediately useful to anyone tasked with the creation of data models for the integration of large-scale enterprise data.

Updated for the latest database management systems -- including MySQL 6.0, Oracle 11g, and Microsoft's SQL Server 2008 -- this introductory guide will get you up and running with SQL quickly. Whether you need to write database applications, perform administrative tasks, or generate reports, Learning SQL, Second Edition, will help you easily master all the SQL fundamentals. Each chapter presents a self-contained lesson on a key SQL concept or technique, with numerous illustrations and annotated examples. Exercises at the end of each chapter let you practice the skills you learn. With this book, you will: Move quickly through SQL basics and learn several advanced features Use SQL data statements to generate, manipulate, and retrieve data Create database objects, such as tables, indexes, and constraints, using SQL schema statements Learn how data sets interact with queries, and understand the importance of subqueries Convert and manipulate data with SQL's built-in functions, and use conditional logic in data statements Knowledge of SQL is a must for interacting with data. With Learning SQL, you'll quickly learn how to put the power and flexibility of this language to work.

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

This guidebook, and its companion volume which follows, provide a solid basis from which one can successfully implement relational database, multidimensional data warehouse and business intelligence (BI) technologies. The principal objective of this initial course volume is to convey a practical and common sense guide to the theory and concepts of data modeling. Using these sophisticated techniques one can create an elegant logical design of a database. Within this course we discuss not only the premier modeling theories from the best industry experts but also present the practical and real-world experience of the past 20-years of Sideris data design practitioners. The methodologies discussed are applicable to any relational database environment, including IBM DB2, the Oracle database, Microsoft SQL Server, the open-source MySQL and PostgreSQL databases as well as other RDBMS platforms. They are also applicable to other database technologies, such as object databases and legacy IMS and IDMS databases. Finally, while we use the free Oracle SQL Developer Data Modeler product as a demonstration modeling tool, one can complete the exercises of this course and apply the techniques learned using any other popular data model diagramming tool, such as IBM InfoSphere Data Architect, CA ErWin Data Modeler, Embarcadero ER/Studio and others. A summary of the objectives of this textbook are: DATA MODELING THEORY & CONCEPTS; BUILDING AN INITIAL DATA MODEL; DRAWING A MODEL USING SOFTWARE ENGINEERING TOOLS; INCREASING THE ACCURACY OF THE MODEL; FINDING & FIXING ATTRIBUTE MISTAKES; SEMANTIC & OBJECT ORIENTED MODELING OF ENTITIES & RELATIONSHIPS; SEMANTIC & OBJECT ORIENTED MODELING OF DOMAINS & TYPES; TIME-DEPENDENCY & STATE-DEPENDENCY; CLASSIC STRUCTURES & PATTERNS; LOGICAL / PHYSICAL MODEL TRANSFORMATION; RDBMS IMPLEMENTATION OF THE PHYSICAL MODEL Database Modeling and Design, Fourth Edition, the extensively revised edition of the classic logical database design reference, explains how you can model and design your database application in consideration of new technology or new business needs. It is an ideal text for a stand-alone data management course focused on logical database design, or a supplement to an introductory text for introductory database management. This book features clear explanations, lots of terrific examples and an illustrative case, and practical advice, with design rules that are applicable to any SQL-based

system. The common examples are based on real-life experiences and have been thoroughly class-tested. The text takes a detailed look at the Unified Modeling Language (UML-2) as well as the entity-relationship (ER) approach for data requirements specification and conceptual modeling - complemented with examples for both approaches. It also discusses the use of data modeling concepts in logical database design; the transformation of the conceptual model to the relational model and to SQL syntax; the fundamentals of database normalization through the fifth normal form; and the major issues in business intelligence such as data warehousing, OLAP for decision support systems, and data mining. There are examples for how to use the most popular CASE tools to handle complex data modeling problems, along with exercises that test understanding of all material, plus solutions for many exercises. Lecture notes and a solutions manual are also available. This edition will appeal to professional data modelers and database design professionals, including database application designers, and database administrators (DBAs); new/novice data management professionals, such as those working on object oriented database design; and students in second courses in database focusing on design. + a detailed look at the Unified Modeling Language (UML-2) as well as the entity-relationship (ER) approach for data requirements specification and conceptual modeling--with examples throughout the book in both approaches! + the details and examples of how to use data modeling concepts in logical database design, and the transformation of the conceptual model to the relational model and to SQL syntax; + the fundamentals of database normalization through the fifth normal form; + practical coverage of the major issues in business intelligence--data warehousing, OLAP for decision support systems, and data mining; + examples for how to use the most popular CASE tools to handle complex data modeling problems. + Exercises that test understanding of all material, plus solutions for many exercises.

The ultimate guide to designing with Oracle8's Object-Relational Model. The authors show users how to implement the concepts in the real world--teaching how to fully exploit the Object-oriented capabilities of Oracle8. They cover the often neglected areas of database design system requirements, like changes to records, data entry errors, and basic transaction history--all key topics that every database designer must address.

Database Modeling and DesignMorgan Kaufmann

A quick and reliable way to build proven databases for core business functions Industry experts raved about The Data Model Resource Book when it was first published in March 1997 because it provided a simple, cost-effective way to design databases for core business functions. Len Silverston has now revised and updated the hugely successful 1st Edition, while adding a companion volume to take care of more specific requirements of different businesses. This updated volume provides a common set of data models for specific core functions shared by most businesses like human resources management, accounting, and project management. These models are standardized and are easily replicated by developers looking for ways to make corporate database development more efficient and cost effective. This guide is the perfect complement to The Data Model Resource CD-ROM, which is sold separately and provides the powerful design templates discussed in the book in a ready-to-use electronic format. A free demonstration CD-ROM is available with each copy of the print book to allow you to try before you buy the full CD-ROM.

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.

Master the most agile and resilient design for building analytics applications: the Unified Star Schema (USS) approach. The USS has many benefits over traditional dimensional modeling. Witness the power of the USS as a single star schema that serves as a foundation for all present and future business requirements of your organization. Data warehouse legend Bill Inmon and business intelligence innovator, Francesco Puppini, explain step-by-step why the Unified Star Schema is the recommended approach for business intelligence designs today, and show through many examples how to build and use this new solution. This book contains two parts. Part I, Architecture, explains the benefits of data marts and data warehouses, covering how organizations progressed to their current state of analytics, and to the challenges that result from current business intelligence architectures. Chapter 1 covers the drivers behind and the characteristics of the data warehouse and data mart. Chapter 2 introduces dimensional modeling concepts, including fact tables, dimensions, star joins, and snowflakes. Chapter 3 recalls the evolution of the data mart. Chapter 4 explains Extract, Transform, and Load (ETL), and the value ETL brings to reporting. Chapter 5 explores the Integrated Data Mart Approach, and Chapter 6 explains how to monitor this environment. Chapter 7 describes the different types of metadata within the data warehouse environment. Chapter 8 progresses through the evolution to our current modern data warehouse environment. Part II, the Unified Star Schema, covers the Unified Star Schema (USS) approach and how it solves the challenges introduced in Part I. There are eight chapters within Part II: · Chapter 9, Introduction to the Unified Star Schema: Learn about its architecture and use cases, as well as how the USS approach differs from the traditional approach. · Chapter 10, Loss of Data: Learn about the loss of data and the USS Bridge. Understand that the USS approach does not create any join, and for this reason, it has no loss of data. · Chapter 11, The Fan Trap: Get introduced to the Oriented Data Model convention, and learn the dangers of a fan trap through an example. Differentiate join and association, and realize that an "in-memory association" is the preferred solution to the fan trap. · Chapter 12, The Chasm Trap: Become familiar with the Cartesian product, and then follow along with an example based on LinkedIn, which illustrates that a chasm trap produces unwanted duplicates. See that the USS Bridge is based on a union, which does not create any duplicates. · Chapter 13, Multi-Fact Queries: Distinguish between multiple facts "with direct connection" versus multiple facts "with no direct connection". Explore how BI tools are capable of building aggregated virtual rows. · Chapter 14, Loops: Learn more about loops and five traditional techniques to solve them. Follow along with an implementation, which will illustrate the solution based on the USS approach. · Chapter 15, Non-Conformed Granularities: Learn about non-conformed granularities, and learn that the Unified Star Schema introduces a solution called "re-normalization". · Chapter 16, Northwind Case Study. Witness how easy it is to detect the pitfalls of Northwind using the ODM convention. Follow along with an implementation of the USS approach on the Northwind database with various BI tools.

Get complete instructions for manipulating, processing, cleaning, and crunching datasets in Python. Updated for Python 3.6, the second edition of this hands-on guide is packed with practical case studies that show you how to solve a broad set of data analysis problems effectively. You'll learn the latest versions of pandas, NumPy, IPython, and Jupyter in the process. Written by Wes McKinney, the creator of the Python pandas project, this book is a practical, modern introduction to data science tools in Python. It's ideal for analysts new to Python and for Python programmers new to data science and scientific computing. Data files and related material are available on GitHub. Use the IPython shell and Jupyter notebook for exploratory computing Learn basic and advanced features in NumPy (Numerical Python) Get started with data analysis tools in the pandas library Use flexible tools to load, clean, transform, merge, and reshape data Create informative visualizations with matplotlib Apply the pandas groupby facility to slice, dice, and summarize datasets Analyze and manipulate regular and irregular time series data Learn how to solve real-world data analysis problems with thorough, detailed examples

This new book aims to provide both beginners and experts with a completely algorithmic approach to data analysis and conceptual modeling, database design, implementation, and tuning, starting from vague and incomplete customer requests and ending with IBM DB/2, Oracle, MySQL, MS SQL Server, or Access based software applications. A rich panoply of solutions to actual useful data sub-universes (e.g. business, university, public and home library, geography, history, etc.) is provided, constituting a powerful library of examples. Four data models are presented and used: the graphical Entity-Relationship, the mathematical EMDM, the physical Relational, and the logical deterministic deductive Datalog ones. For each one of them, best practice rules, algorithms, and the theory beneath are clearly separated. Four

case studies, from a simple public library example, to a complex geographical study are fully presented, on all needed levels. Several dozens of real-life exercises are proposed, out of which at least one per chapter is completely solved. Both major historical and up-to-date references are provided for each of the four data models considered. The book provides a library of useful solutions to real-life problems and provides valuable knowledge on data analysis and modeling, database design, implementation, and fine tuning.

Craft the Right Design Using UML Whether building a relational, object-relational, or object-oriented database, database developers are increasingly relying on an object-oriented design approach as the best way to meet user needs and performance criteria. This book teaches you how to use the Unified Modeling Language-the official standard of the Object Management Group-to develop and implement the best possible design for your database. Inside, the author leads you step by step through the design process, from requirements analysis to schema generation. You'll learn to express stakeholder needs in UML use cases and actor diagrams, to translate UML entities into database components, and to transform the resulting design into relational, object-relational, and object-oriented schemas for all major DBMS products. Features Teaches you everything you need to know to design, build, and test databases using an OO model. Shows you how to use UML, the accepted standard for database design according to OO principles. Explains how to transform your design into a conceptual schema for relational, object-relational, and object-oriented DBMSs. Offers practical examples of design for Oracle, SQL Server, Sybase, Informix, Object Design, POET, and other database management systems. Focuses heavily on re-using design patterns for maximum productivity and teaches you how to certify completed designs for re-use.

This new book refines, customizes, and extends the general Object Modeling Technique (OMT) methodology for the specific subject matter of database applications. By restricting the scope of coverage, the authors are able to present more focused examples and elaborate upon the appropriate methodological steps. The authors present a uniform treatment that addresses files, relational databases, and object-oriented databases.

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.

"This book includes an introduction to fuzzy logic, fuzzy databases and an overview of the state of the art in fuzzy modeling in databases"--Provided by publisher.

This book is a comprehensive presentation of entity-relationship (ER) modeling with regard to an integrated development and modeling of database applications. It comprehensively surveys the achievements of research in this field and deals with the ER model and its extensions. In addition, the book presents techniques for the translation of the ER model into classical database models and languages, such as relational, hierarchical, and network models and languages, as well as into object-oriented models.

This book brings all of the elements of database design together in a single volume, saving the reader the time and expense of making multiple purchases. It consolidates both introductory and advanced topics, thereby covering the gamut of database design methodology ? from ER and UML techniques, to conceptual data modeling and table transformation, to storing XML and querying moving objects databases. The proposed book expertly combines the finest database design material from the Morgan Kaufmann portfolio. Individual chapters are derived from a select group of MK books authored by the best and brightest in the field. These chapters are combined into one comprehensive volume in a way that allows it to be used as a reference work for those interested in new and developing aspects of database design. This book represents a quick and efficient way to unite valuable content from leading database design experts, thereby creating a definitive, one-stop-shopping opportunity for customers to receive the information they would otherwise need to round up from separate sources. Chapters contributed by various recognized experts in the field let the reader remain up to date and fully informed from multiple viewpoints. Details multiple relational models and modeling languages, enhancing the reader's technical expertise and familiarity with design-related requirements specification. Coverage of both theory and practice brings all of the elements of database design together in a single volume, saving the reader the time and expense of making multiple purchases.

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.

With the aim of simplifying relational database modeling, Database Modeling Step-by-Step presents the standard approach to database normalization and then adds its own approach, which is a more simplistic, intuitive way to building relational database models. Going from basics to contemporary topics, the book opens with relational data modeling and ends with BigData

database modeling following a road map of the evolution in relational modeling and including brief introductions to data warehousing and BigData modeling. A break-down of the elements of a model explains what makes up a relational data model. This is followed by a comparison between standard normalization and a more simplistic intuitive approach to data modeling that a beginner can follow and understand. A brief chapter explains how to use the database programming language SQL (Structured Query Language), which reads from and writes to a relational database. SQL is fundamental to data modeling because it helps in understanding how the model is used. In addition to the relational model, the last three chapters cover important modern world topics including denormalization that leads into data warehouses and BigData database modeling. The book explains how there is not much to logical data modeling in BigData databases because as they are often schema-less, which means that BigData databases do not have schemas embedded into the database itself, they have no metadata and thus not much of a logical data model. Online bonus chapters include a case study that covers relational data modeling and are available at the author's web site:

www.oracletroubleshooter.com/datamodeling.html

The purpose of this book is to provide a practical approach for IT professionals to acquire the necessary knowledge and expertise in data modeling to function effectively. It begins with an overview of basic data modeling concepts, introduces the methods and techniques, provides a comprehensive case study to present the details of the data model components, covers the implementation of the data model with emphasis on quality components, and concludes with a presentation of a realistic approach to data modeling. It clearly describes how a generic data model is created to represent truly the enterprise information requirements.

"This book presents cutting-edge research and analysis of the most recent advancements in the fields of database systems and software development"--Provided by publisher.

This book covers the broad field of database design from the perspective of semantic modeling. Aimed at present and future designers of database applications, software engineers, systems analysts and programmers, it aims to offer a unified study of semantic, relational, network and hierarchical databases as seen through the semantic modeling approach. The book provides a structured top-down methodology of database design in all the models and presents the principal types of database languages.

Did you ever try getting Businesspeople and IT to agree on the project scope for a new application? Or try getting Marketing and Sales to agree on the target audience? Or try bringing new team members up to speed on the hundreds of tables in your data warehouse — without them dozing off? Whether you are a businessperson or an IT professional, you can be the hero in each of these and hundreds of other scenarios by building a High-Level Data Model. The High-Level Data Model is a simplified view of our complex environment. It can be a powerful communication tool of the key concepts within our application development projects, business intelligence and master data management programs, and all enterprise and industry initiatives. Learn about the High-Level Data Model and master the techniques for building one, including a comprehensive ten-step approach and hands-on exercises to help you practice topics on your own. In this book, we review data modeling basics and explain why the core concepts stored in a high-level data model can have significant business impact on an organization. We explain the technical notation used for a data model and walk through some simple examples of building a high-level data model. We also describe how data models relate to other key initiatives you may have heard of or may be implementing in your organization. This book contains best practices for implementing a high-level data model, along with some easy-to-use templates and guidelines for a step-by-step approach. Each step will be illustrated using many examples based on actual projects we have worked on. Names have been changed to protect the innocent, but the pain points and lessons have been preserved. One example spans an entire chapter and will allow you to practice building a high-level data model from beginning to end, and then compare your results to ours. Building a high-level data model following the ten step approach you'll read about is a great way to ensure you will retain the new skills you learn in this book. As is the case in many disciplines, using the right tool for the right job is critical to the overall success of your high-level data model implementation. To help you in your tool selection process, there are several chapters dedicated to discussing what to look for in a high-level data modeling tool and a framework for choosing a data modeling tool, in general. This book concludes with a real-world case study that shows how an international energy company successfully used a high-level data model to streamline their information management practices and increase communication throughout the organization—between both businesspeople and IT. Data modeling is one of the under-exploited, and potentially very valuable, business capabilities that are often hidden away in an organization's Information Technology department. Data Modeling for the Business highlights both the resulting damage to business value, and the opportunities to make things better. As an easy-to follow and comprehensive guide on the 'why' and 'how' of data modeling, it also reminds us that a successful strategy for exploiting IT depends at least as much on the information as the technology. Chris Potts, Corporate IT Strategist and Author of *fruitTion: Creating the Ultimate Corporate Strategy for Information Technology* One of the most critical systems issues is aligning business with IT and fulfilling business needs using data models. The authors of *Data Modeling for the Business* do a masterful job at simply and clearly describing the art of using data models to communicate with business representatives and meet business needs. The book provides many valuable tools, analogies, and step-by-step methods for effective data modeling and is an important contribution in bridging the much needed connection between data modeling and realizing business requirements. Len Silverston, author of *The Data Model Resource Book* series

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

Best-selling author and database expert with more than 25 years of experience modeling application and enterprise data, Dr. Michael Blaha provides tried and tested data model patterns, to help readers avoid common modeling mistakes and unnecessary frustration on their way to building effective data models. Unlike the typical methodology book, Patterns of Data Modeling provides advanced techniques for those who have mastered the basics. Recognizing that database representation sets the path for software, determines its flexibility, affects its quality, and influences whether it succeeds or fails, the text focuses on databases rather than programming. It is one of the first books to apply the popular patterns perspective to database systems and data models. It offers practical advice on the core aspects of applications and provides authoritative coverage of mathematical templates, antipatterns, archetypes, identity, canonical models, and relational database design.

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.

This work has been revised and updated to provide a comprehensive treatment of database design for commercial database products and their applications. The book covers the basic foundation of design as well as more advanced techniques, and also incorporates coverage of data warehousing and OLAP (On-Line Analytical Processing), data mining, object-relational, multimedia, and temporal/spatial design.

This is the digital version of the printed book (Copyright © 1996). Learning the basics of a modeling technique is not the same as learning how to use and apply it. To develop a data model of an organization is to gain insights into its nature that do not come easily. Indeed, analysts are often expected to understand subtleties of an organization's structure that may have evaded people who have worked there for years. Here's help for those analysts who have learned the basics of data modeling (or "entity/relationship modeling") but who need to obtain the insights required to prepare a good model of a real business. Structures common to many types of business are analyzed in areas such as accounting, material requirements planning, process manufacturing, contracts, laboratories, and documents. In each chapter, high-level data models are drawn from the following business areas: The Enterprise and Its World The Things of the Enterprise Procedures and Activities Contracts Accounting The Laboratory Material Requirements Planning Process Manufacturing Documents Lower-Level Conventions Typically, analysis, development, and database teams work for different business units, and use different design notations. With UML and the Rational Unified Process (RUP), however, they can unify their efforts -- eliminating time-consuming, error-prone translations, and accelerating software to market. In this book, two data modeling specialists from Rational Software Corporation show exactly how to model data with UML and RUP, presenting proven processes and start-to-finish case studies. The book utilizes a running case study to bring together the entire process of data modeling with UML. Each chapter dissects a different stage of the data modeling process, from requirements through implementation. For each stage, the authors cover workflow and participants' roles, key concepts, proven approach, practical design techniques, and more. Along the way, the authors demonstrate how integrating data modeling into a unified software design process not only saves time and money, but gives all team members a far clearer understanding of the impact of potential changes. The book includes a detailed glossary, as well as appendices that present essential Use Case Models and descriptions. For all software team members: managers, team leaders, systems and data analysts, architects, developers, database designers, and others involved in building database applications for the enterprise.

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.

[Copyright: 802bcb5570392bf25b19d01326b769f](#)