

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

Choose the right Azure data service and correct model design for successful implementation 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.

Imagine what you could do if scalability wasn't a problem. With this hands-on guide, you'll learn how the Cassandra database management system handles hundreds of terabytes of data while remaining highly available across multiple data centers. This expanded second edition—updated for Cassandra 3.0—provides the technical details and practical examples you need to put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's non-relational design, with special attention to data modeling. If you're a developer, DBA, or application architect looking to solve a database scaling issue or future-proof your application, this guide helps you harness Cassandra's speed and flexibility. Understand Cassandra's distributed and decentralized structure Use the Cassandra Query Language (CQL) and cqlsh—the CQL shell Create a working data model and compare it with an equivalent relational model Develop sample applications using client drivers for languages including Java, Python, and Node.js Explore cluster topology and learn how nodes exchange data Maintain a high level of performance in your cluster Deploy Cassandra on site, in the Cloud, or with Docker Integrate Cassandra with Spark, Hadoop, Elasticsearch, Solr, and Lucene

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.

Data modeling is one of the most critical phases in the database application development process, but also the phase most likely to fail. A master data modeler must come into any organization, understand its data requirements, and skillfully model the data for applications that most effectively serve organizational needs. Mastering Data Modeling is a complete guide to becoming a successful data modeler. Featuring a requirements-driven approach, this book clearly explains fundamental concepts, introduces a user-oriented data modeling notation, and describes a rigorous, step-by-step process for collecting, modeling, and documenting the kind of data that users need. Assuming no prior knowledge, Mastering Data Modeling sets forth several fundamental problems of data modeling, such as reconciling the software developer's demand for rigor with the users' equally valid need to speak their own (sometimes vague) natural language. In addition, it describes the good habits that help you respond to these fundamental problems. With these good habits in mind, the book describes the Logical Data Structure (LDS) notation and the process of controlled evolution by which you can create low cost, user-approved data models that resist premature obsolescence. Also included is an encyclopedic analysis of all data shapes that you will encounter. Most notably, the book describes The Flow, a loosely scripted process by which you and the users gradually but continuously improve an LDS until it faithfully represents the information needs. Essential implementation and technology issues are also covered. You will learn about such vital topics as: The fundamental problems of data modeling The good habits that help a data modeler be effective and economical LDS notation, which encourages these good habits How to read an LDS aloud--in declarative English sentences How to write a well-formed (syntactically correct) LDS How to get users to name the parts of an LDS with words from their own business vocabulary How to visualize data for an LDS A catalog of LDS shapes that recur throughout a data models The Flow--the template for your conversations with users How to document an LDS for users, data modelers, and technologists How to map an LDS to a relational schema How LDS differs from other notations and why "Story interludes" appear throughout the book illustrating real-world successes of the LDS notation and controlled evolution process. Numerous exercises help you master critical skills. In addition, two detailed, annotated sample conversations with users show you the process of controlled evolution in action.

Data Modeling Theory and Practice

Data Modeling, A Beginner's Guide

Developing High Quality Data Models

A Practical Guide to Database Design

Handbook of Relational Database Design

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.*

*Database Modeling and Design Morgan Kaufmann*

*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 Oppel 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*.

With our appetites for data on the rise, it has become more important than ever to use UML (Unified Modeling Language) to capture and precisely represent all of these data requirements. Learn how to construct UML data models by working through a series of exercises and self-assessment tests. Beginners can learn the UML directly. Experienced modelers can leverage their understanding of existing database notations, as the book extensively compares the UML to traditional data modeling (Information Engineering). 1. Discover a new way of representing data requirements and communicating better with your business customers. 2. Understand what UML constructs mean and how to properly use them. 3. Learn subtleties of the UML. Become a power UML developer. 4. Practice constructing data models with the exercises. The back of the book answers every exercise. 5. Assess your mastery of the material. Each part has a multiple-choice test that can quantify your understanding. 6. Improve your ability to abstract – think about different ways of representation – as you construct data models. 7. Measure the quality of your data models. 8. Be able to create database designs (DDL code) starting from a UML data model. 9. Be able to write SQL database queries using a data model as a blueprint. 10. Know the differences among operational models, data warehouse models, enterprise models, and master models. They are all aspects of data modeling. This book is concise and to the point. You will learn by induction through reading, practice, and feedback.

*Object-oriented Modeling and Design for Database Applications*

*NoSQL and SQL Data Modeling*

*The Data Warehouse Toolkit*

*Mastering Data Modeling*

*The Complete Guide to Dimensional Modeling*

*Using UML for Data Modeling*

***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.*

*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***

*The only Oracle Press guide to MySQL Workbench explains how to design and model MySQL databases. MySQL Workbench Data Modeling and Development helps developers learn how to effectively use this powerful product for database modeling, reverse engineering, and interaction with the database without writing SQL statements. MySQL Workbench is a graphical user interface that can be used to create and maintain MySQL databases without coding. The book covers the interface and explains how to accomplish each step by illustrating best practices visually. Clear examples, instructions, and explanations reveal, in a single volume, the art of database modeling. This Oracle Press guide shows you how to get the tool to do what you want. Annotated screen shots demonstrate all interactions with the tool, and text explains the how, what, and why of each step. Complete coverage Installation and Configuration; Creating and Managing Connections; Data Modeling Concepts; Creating an ERD; Defining the Physical Schemata; Creating and Managing Tables; Creating and Managing Relationships; Creating and Managing Views; Creating and Managing Routines; Creating and Managing Routine Groups; Creating and Managing User & Groups; Creating and Managing SQL Scripts; Generating SQL Scripts; Forward Engineering a*

***Data Model; Synchronize a Model with a Database; Reverse Engineering a Database; Managing Differences in the Data Catalog; Creating and Managing Model Notes; Editing Table Data; Editing Generated Scripts; Creating New Instances; Managing Import and Export; Managing Security; Managing Server Instances***

***Written from a software engineering perspective, this book shows programmers & developers how to build object-oriented database applications for distributed & client/server environments using the newest update of the OMT methodology & UML.***

***Applied Practices and Principles for Production Ready Software Development***

***Agile Data Warehouse Design***

***Bringing Together Data, Semantics, and Software***

***Patterns of Data Modeling***

***From Conceptual Analysis to Logical Design***

From the #1 source for computing information, trusted by more than six million readers worldwide.

Written for database novices and pros alike, this new book helps you strengthen your command of the underlying theory and practical techniques associated with data modeling and design -- and shows you how to better apply these concepts to today's newest database architectures.

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

Six-Step Relational Database Design™ bridges the gaps between database theory, database modeling, and database implementation by outlining a simple but reliable six-step process for accurately modeling user data on a Crow's Foot Relational Model Diagram, and then demonstrating how to implement this model on any relational database management system. The second edition contains a new chapter on implementation that goes through the steps necessary to implement each of the case studies on a relational database management system, clearly relating the design to implementation and database theory. In addition, questions are also included at the end of each of the six steps and one of the previous case studies has been replaced, making the case study selection more diverse. Six-Step Relational Database Design™ uses three case studies and starts with a statement of the problem by the client and then goes through the six steps necessary to create a reliable and accurate data model of the client's business requirements. This model can then be used to implement the database on any relational database management system. Six-Step Relational Database Design™ should be used as a handbook for students and professionals in the software-development field. The technique described in this book can be used by students for quickly developing relational databases for their applications, and by professionals for developing sturdy, reliable, and accurate relational database models for their software applications.

Implement professional data design and structures in Azure

A User Driven Approach

Design - Build - Run

Fuzzy Databases

Oracle8 Design Using UML Object Modeling

UML Database Modeling Workbook

"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.

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. 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.

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.

Effective database design techniques for data architects and business intelligence professionals

Logical Design

The Language of SQL

Database Modeling and Design - The Entity-Relationship Approach  
Modeling, Design, and Implementation

Information Modeling and Relational Databases

***Fully updated and expanded from the previous edition, A Practical Guide to***

**Database Design, Second Edition, is intended for those involved in the design or development of a database system or application. It begins by focusing on how to create a logical data model where data is stored "where it belongs." Next, data usage is reviewed to transform the logical model into a physical data model that will satisfy user performance requirements. Finally, it describes how to use various software tools to create user interfaces to review and update data in a database. Organized into 11 chapters, the book begins with an overview of the functionality of database management systems and how they guarantee the accuracy and availability of data. It then describes how to define and normalize data requirements to create a logical data model, then map them into an initial solution for a physical database. The book next presents how to use an industry-leading data modeling tool to define and manage logical and physical data models. After that, it describes how to implement a physical database using either Microsoft Access or SQL Server and how to use Microsoft Access to create windows interfaces to query or update data in tables. The last part of the book reviews software tools and explores the design and implementation of a database using as an example a much more complex data environment for a University. The book ends with a description of how to use PHP to build a web-based interface to review and update data in a database.**

**Solve all big data problems by learning how to create efficient data models**  
**Key Features**  
**Create effective models that get the most out of big data**  
**Apply your knowledge to datasets from Twitter and weather data to learn big data**  
**Tackle different data modeling challenges with expert techniques presented in this book**  
**Book Description** Modeling and managing data is a central focus of all big data projects. In fact, a database is considered to be effective only if you have a logical and sophisticated data model. This book will help you develop practical skills in modeling your own big data projects and improve the performance of analytical queries for your specific business requirements. To start with, you'll get a quick introduction to big data and understand the different data modeling and data management platforms for big data. Then you'll work with structured and semi-structured data with the help of real-life examples. Once you've got to grips with the basics, you'll use the SQL Developer Data Modeler to create your own data models containing different file types such as CSV, XML, and JSON. You'll also learn to create graph data models and explore data modeling with streaming data using real-world datasets. By the end of this book, you'll be able to design and develop efficient data models for varying data sizes easily and efficiently. What you will learn

**Get insights into big data and discover various data models**  
**Explore conceptual, logical, and big data models**  
**Understand how to model data containing different file types**  
**Run through data modeling with examples of Twitter, Bitcoin, IMDB and weather data modeling**  
**Create data models such as Graph Data and**

**Vector SpaceModel structured and unstructured data using Python and RWho this book is for This book is great for programmers, geologists, biologists, and every professional who deals with spatial data. If you want to learn how to handle GIS, GPS, and remote sensing data, then this book is for you. Basic knowledge of R and QGIS would be helpful.**

**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.**

**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 Design for Smarties**

**Database Modeling and Design**

**Data Modeling Essentials**

**Cassandra: The Definitive Guide**

**Database Modeling Step by Step**

**Hands-On Big Data Modeling**

**Developing High Quality Data Models provides an introduction to the key principles of data modeling. It explains the purpose of data models in both developing an Enterprise Architecture and in supporting Information Quality; common problems in data model development; and how to develop high quality data models, in particular conceptual, integration, and enterprise data models. The book is organized into**

four parts. Part 1 provides an overview of data models and data modeling including the basics of data model notation; types and uses of data models; and the place of data models in enterprise architecture. Part 2 introduces some general principles for data models, including principles for developing ontologically based data models; and applications of the principles for attributes, relationship types, and entity types. Part 3 presents an ontological framework for developing consistent data models. Part 4 provides the full data model that has been in development throughout the book. The model was created using Jotne EPM Technologys EDMVisualExpress data modeling tool. This book was designed for all types of modelers: from those who understand data modeling basics but are just starting to learn about data modeling in practice, through to experienced data modelers seeking to expand their knowledge and skills and solve some of the more challenging problems of data modeling. Uses a number of common data model patterns to explain how to develop data models over a wide scope in a way that is consistent and of high quality Offers generic data model templates that are reusable in many applications and are fundamental for developing more specific templates Develops ideas for creating consistent approaches to high quality data models 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 unique and critical book shares no-fail secrets for building software and offers tried-and-true practices and principles for software design, development, and testing for mission-critical systems that must not fail. A veteran software architect walks you through the lifecycle of a project as well as each area of production readiness—functionality, availability, performance and scalability, operability, maintainability, and extensibility, and highlights their key concepts.*

*This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. The Language of SQL, Second Edition Many SQL texts attempt to serve as an encyclopedic reference on SQL syntax -- an approach that is often counterproductive, because that information is readily available in online references published by the major database vendors. For SQL beginners, it's more important for a book to focus on general concepts and to offer clear explanations and examples of what various SQL statements can accomplish. This is that book. A number of features make The Language of SQL unique among introductory SQL books. First, you will not be required to download software or sit with a computer as you read the text. The intent of this book is to provide examples of SQL usage that can be understood simply by reading. Second, topics are organized in an intuitive and logical sequence. SQL keywords are introduced one at a time, allowing you to grow your understanding as you encounter new terms and concepts. Finally, this book covers the syntax of three widely used databases: Microsoft SQL Server, MySQL, and Oracle. Special "Database Differences" sidebars clearly show you any differences in syntax among these three databases, and instructions are included on how to obtain and install free versions of the databases. This is the only book you need to gain a quick working knowledge of SQL and relational databases. ·Learn How To... Use SQL to retrieve data from relational databases Apply functions and calculations to data Group and summarize data in a variety of useful ways Use complex logic to retrieve only the data you need Update data and create new tables Design relational databases so that data retrieval is easy and intuitive Use spreadsheets to transform your data into meaningful displays Retrieve data from multiple tables via joins, subqueries, views, and set logic Create, modify, and execute stored procedures Install Microsoft SQL Server, MySQL, or Oracle*

*A Best-Practice Approach to Building Quality Data Models*

*The Shortest Advisable Path*

*MySQL Workbench: Data Modeling & Development*

*UML for Database Design*

*From Logical Data Modeling through Physical Schema Definition*

*Usage-Driven Database Design*

*This book provides a practical and proven approach to designing relational databases. It contains two complementary design methodologies: logical data modeling and relational database design. The design methodologies are independent of product-specific implementations and have been applied to numerous relational product environments. 0201114348B04062001*

*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.*

*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.*

*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.*

*Beginning Database Design*

*The Semantic Modeling Approach*

*Data Modeling and Design for Today's Architectures*

*The Entity-relationship Approach*

*Collaborative Dimensional Modeling, from Whiteboard to Star Schema*

*The Data Modeling Handbook*

This practical, field-tested reference doesn't just explain the characteristics of finished, high-quality data models--it shows readers exactly how to build one. It presents rules and best practices in several notations, including IDEFIX, Martin, Chen, and Finkelstein. The book offers dozens of real-world examples and go beyond basic theory to provide users with practical guidance.

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](http://www.oracletroubleshooter.com/datamodeling.html)

Neo4j is a graph database that allows you to model your data as a graph and find solutions to complex real-world problems that are difficult to solve using any other type of database. This book is designed to help you understand the intricacies of modeling a graph for any domain. The book starts with an example of a graph problem and then introduces you to modeling non-graph problems using Neo4j. Concepts such as the evolution of your database, chains, access control, and recommendations are addressed, along with examples and are modeled in a graph. Throughout the book, you will discover design choices and trade-offs, and understand how and when to use them. By the end of the book, you will be able to effectively use Neo4j to model your database for efficiency and flexibility.

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 dat ...

A Step by Step Approach to Relational Database Design and Development  
Database Modeling and Design, 4th Edition

Conceptual Data Modeling and Database Design: A Fully Algorithmic  
Approach, Volume 1

Data Modeling for Azure Data Services

Neo4j Graph Data Modeling

Distributed Data at Web Scale

*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.*

*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 Theory and Practice is for practitioners and academics who have learned the conventions and rules of data modeling and are looking for a deeper understanding of the discipline. The coverage of theory includes a detailed review of the extensive literature on data modeling and logical database design, referencing nearly 500 publications, with a strong focus on their relevance to practice. The practice component incorporates the largest-ever study of data modeling practitioners, involving over 450 participants in interviews, surveys and data modeling tasks. The results challenge many long-held assumptions about data modeling and will be of interest to academics and practitioners alike.*

*Graeme Simsion brings to the book the practical perspective and intellectual clarity that have made his Data Modeling Essentials a classic in the field. He begins with a question about the nature of data modeling (design or description), and uses it to illuminate such issues as the definition of data modeling, its philosophical underpinnings, inputs and deliverables, the necessary behaviors and skills, the role of creativity, product diversity, quality measures, personal styles, and the differences between experts and novices. Data Modeling Theory and Practice is essential reading for anyone involved in data modeling practice, research, or teaching.*

*Data Modeling and Database Design*

*Data Modeling Logical Database Design*

*Six-Step Relational Database Design*