

The Data Warehouse Toolkit The Complete Guide To Dimensional Modeling

DW 2.0: The Architecture for the Next Generation of Data Warehousing is the first book on the new generation of data warehouse architecture, DW 2.0, by the father of the data warehouse. The book describes the future of data warehousing that is technologically possible today, at both an architectural level and technology level. The perspective of the book is from the top down: looking at the overall architecture and then delving into the issues underlying the components. This allows people who are building or using a data warehouse to see what lies ahead and determine what new technology to buy, how to plan extensions to the data warehouse, what can be salvaged from the current system, and how to justify the expense at the most practical level. This book gives experienced data warehouse professionals everything they need in order to implement the new generation DW 2.0. It is designed for professionals in the IT organization, including data architects, DBAs, systems design and development professionals, as well as data warehouse and knowledge management professionals. * First book on the new generation of data warehouse architecture, DW 2.0. * Written by the "father of the data warehouse", Bill Inmon, a columnist and newsletter editor of The Bill Inmon Channel on the Business Intelligence Network. * Long overdue comprehensive coverage of the implementation of technology and tools that enable the new generation of the DW: metadata, temporal data, ETL, unstructured data, and data quality control.

Ralph Kimball's three data warehousing books, The Data Warehouse Toolkit, The Data Warehouse Lifecycle Toolkit, and The Data Webhouse Toolkit, provide you with everything you will need to create, manage, and use your data warehouse. His first book, The Data Warehouse Toolkit, is the definitive guide to building a data warehouse. Kimball uses actual case studies of existing data warehouses developed for specific types of business applications such as retail, manufacturing, banking, insurance, subscriptions and airline reservations. Using the techniques learned in Kimball's first book, The Data Warehouse Lifecycle Toolkit carries them to the larger issues of delivering complete data marts and data warehouses. The book shows you all the practical details involved in planning, designing, developing, deploying, and growing data warehouses. The Data Webhouse Toolkit is a groundbreaking guide which introduces the Webhouse, a powerful new way of capturing valuable information flowing into a Web site and ordering it in ways that are useful to managers, strategic decision-makers, and customers.

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

With this textbook, Vaisman and Zimányi deliver excellent coverage of data warehousing and business intelligence technologies ranging from the most basic principles to recent findings and applications. To this end, their work is structured into three parts. Part I describes "Fundamental Concepts" including conceptual and logical data warehouse design, as well as querying using MDX, DAX and SQL/OLAP. This part also covers data analytics using Power BI and Analysis Services. Part II details "Implementation and Deployment," including physical design, ETL and data warehouse design methodologies. Part III covers "Advanced Topics" and it is almost completely new in this second edition. This part includes chapters with an in-depth coverage of temporal, spatial, and mobility data warehousing. Graph data warehouses are also covered in detail using Neo4j. The last chapter extensively studies big data management and the usage of Hadoop, Spark, distributed, in-memory, columnar, NoSQL and NewSQL database systems, and data lakes in the context of analytical data processing. As a key characteristic of the book, most of the topics are presented and illustrated using application tools. Specifically, a case study based on the well-known Northwind database illustrates how the concepts presented in the book can be implemented using Microsoft Analysis Services and Power BI. All chapters have been revised and updated to the latest versions of the software tools used. KPIs and Dashboards are now also developed using DAX and Power BI, and the chapter on ETL has been expanded with the implementation of ETL processes in PostgreSQL. Review questions and exercises complement each chapter to support comprehensive student learning. Supplemental material to assist instructors using this book as a course text is available online and includes electronic versions of the figures, solutions to all exercises, and a set of slides accompanying each chapter. Overall, students, practitioners and researchers alike will find this book the most comprehensive reference work on data warehouses, with key topics described in a clear and educational style. "I can only invite you to dive into the contents of the book, feeling certain that once you have completed its reading (or maybe, targeted parts of it), you will join me in expressing our gratitude to Alejandro and Esteban, for providing such a comprehensive textbook for the field of data warehousing in the first place, and for keeping it up to date with the recent developments, in this current second edition." From the foreword by Panos Vassiliadis, University of Ioannina, Greece.

The Complete Guide to Dimensional Modeling

Applied Data Science

With SQL ServerÂ 2005 and the Microsoft Business Intelligence Toolset

The Microsoft Data Warehouse Toolkit

The Modern Data Warehouse in Azure

Building the Data Warehouse

Three Volume Set of Ralph Kimball's Toolkit Books

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

The first, step-by-step guide to building Web-enabled data warehouses The Web can be an incredibly rich source of customer data, and right now companies across industry sectors are hustling to get up and running with data warehouses capable of capturing the clickstream data from their Web sites. This allows companies to track exactly where a customer is going, or "clicking to," on their site in order to gain meaningful information about that customer's preferences. Following Ralph Kimball's The Data Webhouse Toolkit (0-471-37680-9) where he provides the blueprint, Clickstream Data Warehousing fills developers in on all the technical details that go into building a Web-enabled data warehouse. The authors review all key architectural and design issues that developers need to masterfully build a Webhouse using examples to illustrate key points. Companion Web site features code examples from the book and links to related Web sites.

This updated and expanded second edition of the The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling provides a user-friendly introduction to the subject Taking a clear structural framework, it guides the reader through the subject's core elements. A flowing writing style combines with the use of illustrations and diagrams throughout the text to ensure the reader understands even the most complex of concepts. This succinct and enlightening overview is a required reading for all those interested in the subject . We hope you find this book useful in shaping your future career & Business.

Brilliant. . . . Lewis has given us a spectacular account of two great men who faced up to uncertainty and the limits of human reason. William Easterly, Wall Street Journal Forty years ago, Israeli psychologists Daniel Kahneman and Amos Tversky wrote a series of breathtakingly original papers that invented the field of behavioral economics. One of the greatest partnerships in the history of science, Kahneman and Tversky's extraordinary friendship incited a revolution in Big Data studies, advanced evidence-based medicine, led to a new approach to government regulation, and made much of Michael Lewis's own work possible. In The Undoing Project, Lewis shows how their Nobel Prize-winning theory of the mind altered our perception of reality.

Relational and Dimensional Techniques

Clickstream Data Warehousing

The Data Warehouse Toolkit, 2nd Edition; The Data Warehouse Lifecycle, 2nd Edition; The Data Warehouse ETL Toolkit

Relentlessly Practical Tools for Data Warehousing and Business Intelligence

Building a Scalable Data Warehouse with Data Vault 2.0

The MicrosoftÂ Data Warehouse Toolkit

The Purple Morrow is destroyed. However, its promise endures in the form of a champion, mankind's only hope against the destruction spawned by an ancient, sinister evil. The Rovers have invaded the Southernlands, forcing the inhabitants deeper south for respite. Ex-Rover commander Kelen arrives in Swill, a defunct desert town, with revenge in mind. There he encounters the Shadow Man who reveals the truth of his past, including the disturbing origin of a mysterious mark on his forearm. Kelen tries to fight destiny, but soon learns the Shadow will not relent until a terrible prophecy is fulfilled. As Marathana quails under the threat of burgeoning darkness, Jeru's clan looks to him to lead. Jeru, however, knows his path lies elsewhere. Leaving everything behind, he braves the Badlands, a hellish desert only the hardest dare go, to find the Storyteller who can awaken the latent spirit of the Papilion within him. Jeru learns much in that place of desolation, including one truth which could turn destiny on its head. Now more than ever, Kelen, the Wolf of the North, threatens everything he holds dear. Jeru must choose: save those he loves, or sacrifice them to save Marathana.

Develop a custom, agile data warehousing and business intelligence architecture Empower your users and drive better decision making across your enterprise with detailed instructions and best practices from an expert developer and trainer. The Data Warehouse Mentor: Practical Data Warehouse and Business Intelligence Insights shows how to plan, design, construct, and administer an integrated end-to-end DW/BI solution. Learn how to choose appropriate components, build an enterprise data model, configure data marts and data warehouses, establish data flow, and mitigate risk. Change management, data governance, and security are also covered in this comprehensive guide. Understand the components of BI and data warehouse systems Establish project goals and implement an effective deployment plan Build accurate logical and physical enterprise data models Gain insight into your company's transactions with data mining Input, cleanse, and normalize data using ETL (Extract, Transform, and Load) techniques Use structured input files to define data requirements Employ top-down, bottom-up, and hybrid design methodologies Handle security and optimize performance using data governance tools Robert Laberge is the founder of several Internet ventures and a principle consultant for the IBM Industry Models and Assets Lab, which has a focus on data warehousing and business intelligence solutions.

Cowritten by Ralph Kimball, the world's leading data warehousing authority Delivers real-world solutions for the most time- and labor-intensive portion of data warehousing-data staging, or the extract, transform, load (ETL) process Delineates best practices for extracting data from scattered sources, removing redundant and inaccurate data, transforming the remaining data into correctly formatted data structures, and then loading the end product into the data warehouse Offers proven time-saving ETL techniques, comprehensive guidance on building dimensional structures, and crucial advice on ensuring data quality This book is also available as part of the Kimball's Data Warehouse Toolkit Classics Box Set (ISBN: 9780470479575) with the following 3 books: The Data Warehouse Toolkit, 2nd Edition (9780471200246) The Data Warehouse Lifecycle Toolkit, 2nd Edition (9780470149775) The Data Warehouse ETL Toolkit (9780764567575)

Over the past 5 years, the concept of big data has matured, data science has grown exponentially, and data architecture has become a standard part of organizational decision-making. Throughout all this change, the basic principles that shape the architecture of data have remained the same. There remains a need for people to take a look at the "bigger picture" and to understand where their data fit into the grand scheme of things. Data Architecture: A Primer for the Data Scientist, Second Edition addresses the larger architectural picture of how big data fits within the existing information infrastructure or data warehousing systems. This is an essential topic not only for data scientists, analysts, and managers but also for researchers and engineers who increasingly need to deal with large and complex sets of data. Until data are gathered and can be placed into an existing framework or architecture, they cannot be used to their full potential. Drawing upon years of practical experience and using numerous examples and case studies from across various industries, the authors seek to explain this larger picture into which big data fits, giving data scientists the necessary context for how pieces of the puzzle should fit together. New case studies include expanded coverage of textual management and analytics New chapters on visualization and big data Discussion of new visualizations of the end-state architecture

A Guide to Modern Databases and the NoSQL Movement

The Data Warehouse ETL Toolkit

A Well-Ordered Thing

With SQL Server 2008 R2 and the Microsoft Business Intelligence Toolset

Understanding Web Services

A Primer for the Data Scientist

The Data Warehouse Mentor: Practical Data Warehouse and Business Intelligence Insights

Build a modern data warehouse on Microsoft's Azure Platform that is flexible, adaptable, and fast—fast to snap together, reconfigure, and fast at delivering results to drive good decision making in your business. Gone are the days when data warehousing projects were lumbering dinosaur-style projects that took forever, drained budgets, and produced business intelligence (BI) just in time to tell you what to do 10 years ago. This book will show you how to assemble a data warehouse solution like a jigsaw puzzle by connecting specific Azure technologies that address your own needs and bring value to your business. You will see how to implement a range of architectural patterns using batches, events, and streams for both data lake technology and SQL databases. You will discover how to manage metadata and automation to accelerate the development of your warehouse while establishing resilience at every level. And you will know how to feed downstream analytic solutions such as Power BI and Azure Analysis Services to empower data-driven decision making that drives your business forward toward a pattern of success. This book teaches you how to employ the Azure platform in a strategy to dramatically improve implementation speed and flexibility of data warehousing systems. You will know how to make correct decisions in design, architecture, and infrastructure such as choosing which type of SQL engine (from at least three options) best meets the needs of your organization. You also will learn about ETL/ELT structure and the vast number of accelerators and patterns that can be used to aid implementation and ensure resilience. Data warehouse developers and architects will find this book a tremendous resource for moving their skills into the future through cloud-based implementations. What You Will LearnChoose the appropriate Azure SQL engine for implementing a given data warehouse Develop smart, reusable ETL/ELT processes that are resilient and easily maintained Automate mundane development tasks through tools such as PowerShell Ensure consistency of data by creating and enforcing data contracts Explore streaming and event-driven architectures for data ingestionCreate advanced staging layers using Azure Data Lake Gen 2 to feed your data warehouse Who This Book Is For Data warehouse or ETL/ELT developers who wish to implement a data warehouse project in the Azure cloud, and developers currently working in on-premise environments who want to move to the cloud, and for developers with Azure experience looking to tighten up their implementation and consolidate their knowledge

The "father of data warehousing" incorporates the latest technologies into his blueprint for integrated decision supportsystems Today's corporate IT and data warehouse managers are required to make a small army of technologies work together to ensure fast and accurate information for business managers. Bill Inmon created the Corporate Information Factory to solve the needs of these managers. Since the First Edition, the design of the factory has grown and changed dramatically. This Second Edition, revised and expanded by 40% with five new chapters, incorporates these changes. This step-by-step guide will enable readers to connect their legacy systems with the data warehouse and deal with a host of new and changing technologies, including Web access mechanisms, e-commerce systems, ERP (Enterprise Resource Planning) systems. The book also looks closely at exploration and data mining servers for analyzing customer behavior and departmental data marts for finance, sales, and marketing.

The Data Warehouse ToolkitThe Complete Guide to Dimensional ModelingJohn Wiley & Sons Incorporated

Cowritten by Ralph Kimball, the world's leading data warehousing authority, whose previous books have sold more than 150,000 copies Delivers real-world solutions for the most time- and labor-intensive portion of data warehousing-data staging, or the extract, transform, load (ETL) process Delineates best practices for extracting data from scattered sources, removing redundant and inaccurate data, transforming the remaining data into correctly formatted data structures, and then loading the end product into the data warehouse Offers proven time-saving ETL techniques, comprehensive guidance on building dimensional structures, and crucial advice on ensuring data quality

Mastering Data Warehouse Design

Practical Techniques for Building Dimensional Data Warehouses

Building the Data Lakehouse

XML, WSDL, SOAP, and UDDI

A Hands-on Guide to Relational Database Design

Data Architecture: A Primer for the Data Scientist

Solutions for Star Schema Performance

This is the first book to provide in-depth coverage of star schema aggregates used in dimensional modeling—from selection and design, to loading and usage, to specific tasks and deliverables for implementation projects Covers the principles of aggregate schema design and the pros and cons of various types of commercial solutions for navigating and building aggregates Discusses how to include aggregates in data warehouse development projects that focus on incremental development, iterative builds, and early data loads

"This book takes the somewhat daunting process of database design and breaks it into completely manageable and understandable components. Mike's approach whilst simple is completely professional, and I can recommend this book to any novice database designer." --Sandra Barker, Lecturer, University of South Australia, Australia "Databases are a critical infrastructure technology for information systems and today's business. Mike Hernandez has written a literate explanation of database technology—a topic that is intricate and often obscure. If you design databases yourself, this book will educate you about pitfalls and show you what to do. If you purchase products that use a database, the book explains the technology so that you can understand what the vendor is doing and assess their products better." --Michael Blaha, consultant and trainer, author of A Manager's Guide to Database Technology "If you told me that Mike Hernandez could improve on the first edition of Database Design for Mere Mortals I wouldn't have believed you, but he did! The second edition is packed with more real-world examples, detailed explanations, and even includes database-design tools on the CD-ROM! This is a must-read for anyone who is even remotely interested in relational database design, from the individual who is called upon occasionally to create a useful tool at work, to the seasoned professional who wants to brush up on the fundamentals. Simply put, if you want to do it right, read this book!" --Matt Greer, Process Control Development, The Dow Chemical Company "Mike's approach to database design is totally common-sense based, yet he's adhered to all the rules of good relational database design. I use Mike's books in my starter database-design class, and I recommend his books to anyone who's interested in learning how to design databases or how to write SQL queries." --Michelle Poollet, President, MVDS, Inc. "Slapping together sophisticated applications with poorly designed data will hurt you just as much now as when Mike wrote his first edition, perhaps even more. Whether you're just getting started developing with data or are a seasoned pro; whether you've read Mike's previous book or this is your first; whether you're happier letting someone else design your data or you love doing it yourself--this is the book for you. Mike's ability to

explain these concepts in a way that's not only clear, but fun, continues to amaze me." --From the Foreword by Ken Getz, MCW Technologies, coauthor ASP.NET Developer's JumpStart "The first edition of Mike Hernandez's book Database Design for Mere Mortals was one of the few books that survived the cut when I moved my office to smaller quarters. The second edition expands and improves on the original in so many ways. It is not only a good, clear read, but contains a remarkable quantity of clear, concise thinking on a very complex subject. It's a must for anyone interested in the subject of database design." --Malcolm C. Rubel, Performance Dynamics Associates "Mike's excellent guide to relational database design deserves a second edition. His book is an essential tool for fledgling Microsoft Access and other desktop database developers, as well as for client/server pros. I recommend it highly to all my readers." --Roger Jennings, author of Special Edition Using Access 2002 "There are no silver bullets! Database technology has advanced dramatically, the newest crop of database servers perform operations faster than anyone could have imagined six years ago, but none of these technological advances will help fix a bad database design, or capture data that you forgot to include! Database Design for Mere Mortals(TM), Second Edition, helps you design your database right in the first place!" --Matt Nunn, Product Manager, SQL Server, Microsoft Corporation "When my brother started his professional career as a developer, I gave him Mike's book to help him understand database concepts and make real-world application of database technology. When I need a refresher on the finer points of database design, this is the book I pick up. I do not think that there is a better testimony to the value of a book than that it gets used. For this reason I have wholeheartedly recommended to my peers and students that they utilize this book in their day-to-day development tasks." --Chris Kunicki, Senior Consultant, OfficeZealot.com "Mike has always had an incredible knack for taking the most complex topics, breaking them down, and explaining them so that anyone can 'get it.' He has honed and polished his first very, very good edition and made it even better. If you're just starting out building database applications, this book is a must-read cover to cover. Expert designers will find Mike's approach fresh and enlightening and a source of great material for training others." --John Viescas, President, Viescas Consulting, Inc., author of Running Microsoft Access 2000 and coauthor of SQL Queries for Mere Mortals "Whether you need to learn about relational database design in general, design a relational database, understand relational database terminology, or learn best practices for implementing a relational database, Database Design for Mere Mortals(TM), Second Edition, is an indispensable book that you'll refer to often. With his many years of real-world experience designing relational databases, Michael shows you how to analyze and improve existing databases, implement keys, define table relationships and business rules, and create data views, resulting in data integrity, uniform access to data, and reduced data-entry errors." --Paul Cornell, Site Editor, MSDN Office Developer Center Sound database design can save hours of development time and ensure functionality and reliability. Database Design for Mere Mortals(TM), Second Edition, is a straightforward, platform-independent tutorial on the basic principles of relational database design. It provides a commonsense design methodology for developing databases that work. Database design expert Michael J. Hernandez has expanded his best-selling first edition, maintaining its hands-on approach and accessibility while updating its coverage and including even more examples and illustrations. This edition features a CD-ROM that includes diagrams of sample databases, as well as design guidelines, documentation forms, and examples of the database design process. This book will give you the knowledge and tools you need to create efficient and effective relational databases.

The Data Vault was invented by Dan Linstedt at the U.S. Department of Defense, and the standard has been successfully applied to data warehousing projects at organizations of different sizes, from small to large-size corporations. Due to its simplified design, which is adapted from nature, the Data Vault 2.0 standard helps prevent typical data warehousing failures. "Building a Scalable Data Warehouse" covers everything one needs to know to create a scalable data warehouse end to end, including a presentation of the Data Vault modeling technique, which provides the foundations to create a technical data warehouse layer. The book discusses how to build the data warehouse incrementally using the agile Data Vault 2.0 methodology. In addition, readers will learn how to create the input layer (the stage layer) and the presentation layer (data mart) of the Data Vault 2.0 architecture including implementation best practices. Drawing upon years of practical experience and using numerous examples and an easy to understand framework, Dan Linstedt and Michael Olschimke discuss: How to load each layer using SQL Server Integration Services (SSIS), including automation of the Data Vault loading processes. Important data warehouse technologies and practices. Data Quality Services (DQS) and Master Data Services (MDS) in the context of the Data Vault architecture. Provides a complete introduction to data warehousing, applications, and the business context so readers can get-up and running fast Explains theoretical concepts and provides hands-on instruction on how to build and implement a data warehouse Demystifies data vault modeling with beginning, intermediate, and advanced techniques Discusses the advantages of the data vault approach over other techniques, also including the latest updates to Data Vault 2.0 and multiple improvements to Data Vault 1.0

". . . one of the definitive books of our industry. If you take the time to read only one professional book, make it this book." -W. H. Inmon One of the most dramatic new developments in database design, the dimensional data warehouse is a powerful database model that significantly enhances managers' ability to quickly analyze large, multidimensional data sets. Written by the leading proponent of this revolutionary new approach, this valuable book/CD toolkit outfits you with all the nuts-and-bolts information you need to design, build, manage, and use dimensional data warehouses for virtually any type of business application, as well as software for querying dimensional data warehouses. Employing many real-life case studies of data warehouses, Ralph Kimball provides clear-cut guidelines on how to model data and design data warehouses to support advanced multidimensional decision support systems. Beginning with the relatively simple example of a data warehouse for a grocery store, he progresses, step-by-step, through an increasingly complex array of business applications in retail, manufacturing, banking, insurance, subscriptions, and airline reservations. By the end of the book, you will have mastered the full range of powerful techniques for creating, controlling, and navigating dimensional business databases that are easy to understand and navigate. On the CD-ROM you'll find: * Software for querying dimensional data warehouses. * Working models of all the databases described in the book.

3 Volume Set

The Big Ideas Behind Reliable, Scalable, and Maintainable Systems

The Data Warehouse Toolkit

Dmitrii Mendeleev and the Shadow of the Periodic Table, Revised Edition

Clean Code

Mastering Data Warehouse Aggregates

A Real Guide from Real Experts on Getting the Job You Want!

Data pipelines are the foundation for success in data analytics. Moving data from numerous diverse sources and transforming it to provide context is the difference between having data and actually gaining value from it. This pocket reference defines data pipelines and explains how they work in today's modern data stack. You'll learn common considerations and key decision points when implementing pipelines, such as batch versus streaming data ingestion and build versus buy. This book addresses the most common decisions made by data professionals and discusses foundational concepts that apply to open source frameworks, commercial products, and homegrown solutions. You'll learn: What a data pipeline is and how it works How data is moved and processed on modern data infrastructure, including cloud platforms Common tools and products used by data engineers to build pipelines How pipelines support analytics and reporting needs Considerations for pipeline maintenance, testing, and alerting

Dmitrii Mendeleev (1834–1907) is a name we recognize, but perhaps only as the creator of the periodic table of elements. Generally, little else has been known about him. A Well-Ordered Thing is an authoritative biography of Mendeleev that draws a multifaceted portrait of his life for the first time. As Michael Gordin reveals, Mendeleev was not only a luminary in the history of science, he was also an astonishingly wide-ranging political and cultural figure. From his attack on Spiritualism to his failed voyage to the Arctic and his near-mythical hot-air balloon trip, this is the story of an extraordinary maverick. The ideals that shaped his work outside science also led Mendeleev to order the elements and, eventually, to engineer one of the most fascinating scientific developments of the nineteenth century. A Well-Ordered Thing is a classic work that tells the story of one of the world's most important minds.

The data lakehouse is the next generation of the data warehouse and data lake, designed to meet today's complex and ever-changing analytics, machine learning, and data science requirements. Learn about the features and architecture of the data lakehouse, along with its powerful analytical infrastructure. Appreciate how the universal common connector blends structured, textual, analog, and IoT data. Maintain the lakehouse for future generations through Data Lakehouse Housekeeping and Data Future-proofing. Know how to incorporate the lakehouse into an existing data governance strategy. Incorporate data catalogs, data lineage tools, and open source software into your architecture to ensure your data scientists, analysts, and end users live happily ever after.

An unparalleled collection of recommended guidelines for data warehousing and business intelligence pioneered by Ralph Kimball and his team of colleagues from the Kimball Group. Recognized and respected throughout the world as the most influential leaders in the data warehousing industry, Ralph Kimball and the Kimball Group have written articles covering more than 250 topics that define the field of data warehousing. For the first time, the Kimball Group's incomparable advice, design tips, and best practices have been gathered in this remarkable collection of articles, which spans a decade of data warehousing innovation. Each group of articles is introduced with original commentaries that explain their role in the overall lifecycle methodology developed by the Kimball Group. These practical, hands-on articles are fully updated to reflect current practices and terminology and cover the complete lifecycle—including project planning, requirements gathering, dimensional modeling, ETL, and business intelligence and analytics. This easily referenced collection is nothing less than vital if you are involved with data warehousing or business intelligence in any capacity.

Lessons Learned for the Data-Driven Business

The Informed Company

Designing Data-Intensive Applications

Wolf's Bane

Database Design for Mere Mortals

Analyzing Data with Power BI and Power Pivot for Excel

All of Nonparametric Statistics

A thorough update to the industry standard for designing, developing, and deploying data warehouse and business intelligence systems The world of data warehousing has changed remarkably since the first edition of The Data Warehouse Lifecycle Toolkit was published in 1998. In that time, the data warehouse industry has reached full maturity and acceptance, hardware and software have made staggering advances, and the techniques promoted in the premiere edition of this book have been adopted by nearly all data warehouse vendors and practitioners. In addition, the term "business intelligence" emerged to reflect the mission of the data warehouse: wrangling the data out of source systems, cleaning it, and delivering it to add value to the business. Ralph Kimball and his colleagues have refined the original set of Lifecycle methods and techniques based on their consulting and training experience. The authors understand first-hand that a data warehousing/business intelligence (DW/BI) system needs to change as fast as its surrounding organization evolves. To that end, they walk you through the detailed steps of designing, developing, and deploying a DW/BI system. You'll learn to create adaptable systems that deliver data and analyses to business users so they can make better business decisions. Discusses application-to-application Internet communication, network standards, major architectural approaches, the role of Web services, and eBXML.

Renowned DAX experts Alberto Ferrari and Marco Russo teach you how to design data models for maximum efficiency and effectiveness. How can you use Excel and Power BI to gain real insights into your information? As you examine your data, how do you write a formula that provides the numbers you need? The answers to both of these questions lie with the data model. This book introduces the basic techniques for shaping data models in Excel and Power BI. It's meant for readers who are new to data modeling as well as for experienced data modelers looking for tips from the experts. If you want to use Power BI or Excel to analyze data, the many real-world examples in this book will help you look at your reports in a different way-like experienced data modelers do. As you'll soon see, with the right data model, the correct answer is always a simple one! By reading this book, you will:

- *Gain an understanding of the basics of data modeling, including tables, relationships, and keys*
- *Familiarize yourself with star schemas, snowflakes, and common modeling techniques*
- *Learn the importance of granularity*
- *Discover how to use multiple fact tables, like sales and purchases, in a complex data model*
- *Manage calendar-related calculations by using date tables*
- *Track historical attributes, like previous addresses of customers or manager assignments*
- *Use snapshots to compute quantity on hand*
- *Work with multiple currencies in the most efficient way*
- *Analyze events that have durations, including overlapping durations*
- *Learn what data model you need to answer your specific business questions*

About This Book

- *For Excel and Power BI users who want to exploit the full power of their favorite tools*
- *For BI professionals seeking new ideas for modeling data*

Focus your efforts on the best opportunities --

Data Pipelines Pocket Reference

Corporate Information Factory

The Data Warehouse Lifecycle Toolkit

A Handbook of Agile Software Craftsmanship

The Kimball Group Reader

Design and Implementation

Collaborative Dimensional Modeling, from Whiteboard to Star Schema

Three books by the bestselling authors on Data Warehousing! The most authoritative guides from the inventor of the technique all for a value price. The Data Warehouse Toolkit, 3rd Edition (9781118530801) Ralph Kimball invented a data warehousing technique called "dimensional modeling" and popularized it in his first Wiley book, The Data Warehouse Toolkit. Since this book was first published in 1996, dimensional modeling has become the most widely accepted technique for data warehouse design. Over the past 10 years, Kimball has improved on his earlier techniques and created many new ones. In this 3rd edition, he will provide a comprehensive collection of all of these techniques, from basic to advanced. The Data Warehouse Lifecycle Toolkit, 2nd Edition (9780470149775) Complete coverage of best practices from data warehouse project inception through on-going program management. Updates industry best practices to be in sync with current recommendations of Kimball Group. Streamlines the lifecycle methodology to be more efficient and user-friendly The Data Warehouse ETL Toolkit (9780764567575) shows data warehouse developers how to effectively manage the ETL (Extract, Transform, Load) phase of the data warehouse development lifecycle. The authors show developers the best methods for extracting data from scattered sources throughout the enterprise, removing obsolete, redundant, and inaccurate data, transforming the remaining data into correctly formatted data structures, and then physically loading them into the data warehouse. This book provides complete coverage of proven, time-saving ETL techniques. It begins with a quick overview of ETL fundamentals and the role of the ETL development team. It then quickly moves into an overview of the ETL data structures, both relational and dimensional. The authors show how to build useful dimensional stuctures, providing practical examples of beginning through advanced techniques.

This new edition enhances, extends, and clarifies the concepts and examples presented in the first edition. Topics have been restructured to coherently develop the data warehouse architecture.

Maybe you're a recent college graduate, looking for a successful start to your career. Or an experienced professional, feeling the need to try something new. Either way, a whole host of opportunities await you-but if you really hope to ace that interview and get the job you want, you'll need the right skills to get ahead. So when you're navigating the complex twists and turns of today's changing job market, let I'll Get That Job! serve as your road map and guide. Featuring advice from real HR professionals, headhunters, and team managers, this essential job-hunting companion will let you know exactly what you need to do to increase your chances, from social media presence to writing a great CV. While shedding light on the many myths and outdated "rules" that may actually bog you down in today's job-seeking experience, I'll Get That Job! serves as a source of motivation and encouragement for modern job hunters. After all, with hard work and the right mind-set, it really is possible for you to get that job you've always wanted-and become the most successful version of yourself along the way!

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.

Building with Speed and Agility on Microsoft 's Cloud Platform

Data Warehouse Systems

The Data Warehouse Toolkit - The Complete Guide to Dimensional Modeling

Seven Databases in Seven Weeks

DW 2.0: The Architecture for the Next Generation of Data Warehousing

I'll Get That Job!

How to Build a Cloud-Based Data Stack to Explore and Understand Data

Learn how to manage a modern data stack and get the most out of data in your organization! Thanks to the emergence of new technologies and the explosion of data in recent years, we need new practices for managing and getting value out of data. In the competitive landscape the "best guess" approach—reading blog posts here and there and patching together data practices without any real visibility—is no longer going to hack it. The Informed Company provides definitive direction on how best to leverage it, including cloud computing, columnar storage, cloud ETL tools, and cloud BI tools. You'll learn how to work with Agile methods and set up processes that's right for your company to use your data as a key weapon for your success. . . . You'll discover best practices at each stage, from querying production databases at a small startup all the way to setting up data marts for different business lines of an enterprise. In their work at Chartio, authors Fowler and David have learned that most businesspeople are almost completely clueless when it comes to data. If they are using resources, those resources are outdated, so they're missing out on the latest cloud technologies and advances in data analytics. This book will firm up your understanding of data and bring you into the present with knowledge and what doesn't. Discover the data stack strategies that are working for today's successful small, medium, and enterprise companies Learn the different Agile stages of data organization, and the right one for your team Learn how to maintain Data Lakes a safe and effective, accessible data storage Gain the knowledge you need to architect Data Warehouses and Data Marts Understand your business's level of data sophistication and the steps you can take to get to "level up" your data The Informed Company is the best resource for anyone who wants to work faster and more nimbly, armed with actionable decision-making data.

A cutting-edge response to Ralph Kimball's challenge to the data warehouse community that answers some tough questions about the effectiveness of the relational approach to data warehousing Written by one of the best-known exponents of the Bill Inmon data warehouse Addresses head-on the tough issues raised by Kimball and explains how to choose the best modeling technique for solving common data warehouse design problems Weighs the pros and cons of relational vs. dimensional modeling techniques Focuses on solving modeling problems, including creating and maintaining keys and modeling calendars, hierarchies, transactions, and data quality

The data warehousing bible updated for the new millennium Updated and expanded to reflect the many technological advances occurring since the previous edition, this latest edition of the data warehousing "bible" provides a comprehensive introduction to the operational data stores, the Corporate Information Factory, exploration warehouses, and Web-enabled warehouses. Written by the father of the data warehouse concept, the book also reviews the unique requirements for supporting e-business and explores how, within the traditional data warehouse can be integrated with new technologies to provide enhanced customer service, sales, and support—both online and offline—including near-line data storage techniques.

Looks at the principles and clean code, includes case studies showcasing the practices of writing clean code, and contains a list of heuristics and "smells" accumulated from the process of writing clean code.

Practical Techniques for Extracting, Cleaning, Conforming, and Delivering Data

Agile Data Warehouse Design

Kimball's Data Warehouse Toolkit Classics

The Undoing Project: A Friendship That Changed Our Minds

The Definitive Guide to Dimensional Modeling

This book has two main goals: to define data science through the work of data scientists and their results, namely data products, while simultaneously providing the reader with relevant lessons learned from applied data science projects at the intersection of academia and industry. As such, it is not a replacement for a classical textbook (i.e., it does not elaborate on fundamentals of methods and principles described elsewhere), but systematically highlights the connection between theory, on the one hand, and its application in specific use cases, on the other. With these goals in mind, the book is divided into three parts: Part I pays tribute to the interdisciplinary nature of data science and provides a common understanding of data science terminology for readers with different backgrounds. These six chapters are geared towards drawing a consistent picture of data science and were predominantly written by the editors themselves. Part II then broadens the spectrum by presenting views and insights from diverse authors – some from academia and some from industry, ranging from financial to health and from manufacturing to e-commerce. Each of these chapters describes a fundamental principle, method or tool in data science by analyzing specific use cases and drawing concrete conclusions from them. The case studies presented, and the methods and tools applied, represent the nuts and bolts of data science. Finally, Part III was again written from the perspective of the editors and summarizes the lessons learned that have been distilled from the case studies in Part II. The section can be viewed as a meta-study on data science across a broad range of domains, viewpoints and fields. Moreover, it provides answers to the question of what the mission-critical factors for success in different data science undertakings are. The book targets professionals as well as students of data science: first, practicing data scientists in industry and academia who want to broaden their scope and expand their knowledge by drawing on the authors' combined experience. Second, decision makers in businesses who face the challenge of creating or implementing a data-driven strategy and who want to learn from success stories spanning a range of industries. Third, students of data science who want to understand both the theoretical and practical aspects of data science, vetted by real-world case studies at the intersection of academia and industry.

This text provides the reader with a single book where they can find accounts of a number of up-to-date issues in nonparametric inference. The book is aimed at Masters or PhD level students in statistics, computer science, and engineering. It is also suitable for researchers who want to get up to speed quickly on modern nonparametric methods. It covers a wide range of topics including the bootstrap, the nonparametric delta method, nonparametric regression, density estimation, orthogonal function methods, minimax estimation, nonparametric confidence sets, and wavelets. The book's dual approach includes a mixture of methodology and theory.

Updated new edition of Ralph Kimball's groundbreaking book on dimensional modeling for data warehousing and business intelligence! The first edition of Ralph Kimball's The Data Warehouse Toolkit introduced the industry to dimensional modeling, and now his books are considered the most authoritative guides in this space. This new third edition is a complete library of updated dimensional modeling techniques, the most comprehensive collection ever. It covers new and enhanced star schema dimensional modeling patterns, adds two new chapters on ETL techniques, includes new and expanded business matrices for 12 case studies, and more. Authored by Ralph Kimball and Margy Ross, known worldwide as educators, consultants, and influential thought leaders in data warehousing and business intelligence. Begins with fundamental design recommendations and progresses through increasingly complex scenarios. Presents unique modeling techniques for business applications such as inventory management, procurement, invoicing, accounting, customer relationship management, big data analytics, and more. Draws real-world case studies from a variety of industries, including retail sales, financial services, telecommunications, education, health care, insurance, e-commerce, and more. Design dimensional databases that are easy to understand and provide fast query response with The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling, 3rd Edition.