

The Data Warehouse Toolkit

The Complete Guide To Dimensional Modeling

Yeah, reviewing a books **The Data Warehouse Toolkit The Complete Guide To Dimensional Modeling** could accumulate your close links listings. This is just one of the solutions for you to be successful. As understood, triumph does not recommend that you have fabulous points.

Comprehending as without difficulty as concurrence even more than additional will find the money for each success. bordering to, the notice as well as keenness of this The Data Warehouse Toolkit The Complete Guide To Dimensional Modeling can be taken as without difficulty as picked to act.

Data Warehouse Systems

Alejandro Vaisman

2022-08-16 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.

Data Pipelines Pocket

Reference James Densmore

2021-02-10 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

Mastering Data Warehouse

Aggregates Christopher

Adamson 2012-06-27 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

Kimball's Data

Warehouse Toolkit

Classics Ralph Kimball
2009-04-06 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)
Agile Data Warehouse Design Lawrence Corr
2011-11 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.

Designing Data-Intensive Applications

Martin Kleppmann 2017-03-16 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

Data Architecture: A Primer for the Data Scientist W.H. Inmon

2019-04-30 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

The Data Warehouse Toolkit - The Complete Guide to Dimensional Modeling Raiph Kimball 2004

Professional Hadoop Solutions Boris Lublinsky 2013-09-12 The go-to guidebook for deploying Big Data solutions with Hadoop Today's enterprise architects need to understand how the Hadoop frameworks and APIs fit together, and how they can be integrated to deliver real-world solutions. This book is a practical, detailed guide to building and implementing those solutions, with code-level instruction in the popular Wrox tradition. It covers storing data with HDFS and Hbase, processing data with MapReduce, and automating data processing with Oozie. Hadoop security, running Hadoop with Amazon Web Services, best practices, and automating Hadoop processes in real time are also covered in

depth. With in-depth code examples in Java and XML and the latest on recent additions to the Hadoop ecosystem, this complete resource also covers the use of APIs, exposing their inner workings and allowing architects and developers to better leverage and customize them. The ultimate guide for developers, designers, and architects who need to build and deploy Hadoop applications Covers storing and processing data with various technologies, automating data processing, Hadoop security, and delivering real-time solutions Includes detailed, real-world examples and code-level guidelines Explains when, why, and how to use these tools effectively Written by a team of Hadoop experts in the programmer-to-programmer Wrox style Professional Hadoop Solutions is the reference enterprise architects and developers need to

maximize the power of Hadoop.

Building the Data

Lakehouse Bill Inmon

2021-10 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.

The Microsoft Data Warehouse Toolkit

Joy Mundy 2011-03-08 Focus your efforts on the best opportunities --

The Data Warehouse Toolkit

Ralph Kimball 2011-08-08

DW 2.0: The Architecture for the Next Generation of Data

Warehousing W.H. Inmon

2010-07-28 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.

Building a Scalable Data Warehouse with Data Vault 2.0

Dan Linstedt
2015-09-15 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

The Data Warehouse Toolkit Spencer M. Riley 2015-08-19 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.

Three Volume Set of Ralph Kimball's Toolkit Books

Ralph Kimball 2000-06-01
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.

The Data Warehouse Toolkit

Ralph Kimball
2013-07-01 Updated new edition of Ralph Kimball's groundbreaking book on dimensional modeling for data warehousing and businessintelligence! The first edition of Ralph Kimball's The Data WarehouseToolkit

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.

Testing the Data Warehouse Practicum

Doug Vučević 2012-08 The quality of a data warehouse (DWH) is the elusive aspect of it, not because it is hard to achieve [once we agree what it is], but because it is difficult to describe. We propose the notion that quality is not an attribute or a feature that a product has to possess, but rather a relationship between that product and each and every stakeholder. More

specifically, the relationship between the software quality and the organization that produces the products is explored. Quality of data that populates the DWH is the main concern of the book, therefore we propose a definition for data quality as: "fitness to serve each and every purpose". Methods are proposed throughout the book to help readers achieve data warehouse quality.

Learning to Love Data Science

Mike Barlow
2015-10-27 Until recently, many people thought big data was a passing fad. "Data science" was an enigmatic term. Today, big data is taken seriously, and data science is considered downright sexy. With this anthology of reports from award-winning journalist Mike Barlow, you'll appreciate how data science is fundamentally altering our world, for better and for worse. Barlow paints a picture of the emerging data space in broad strokes.

From new techniques and tools to the use of data for social good, you'll find out how far data science reaches. With this anthology, you'll learn how: Analysts can now get results from their data queries in near real time Indie manufacturers are blurring the lines between hardware and software Companies try to balance their desire for rapid innovation with the need to tighten data security Advanced analytics and low-cost sensors are transforming equipment maintenance from a cost center to a profit center CIOs have gradually evolved from order takers to business innovators New analytics tools let businesses go beyond data analysis and straight to decision-making Mike Barlow is an award-winning journalist, author, and communications strategy consultant. Since launching his own firm, Cumulus Partners, he has represented major

organizations in a number of industries.

Learn Data Warehousing

in 1 Day Krishna Rungta
2018-02-15 Unlike popular belief, Data Warehouse is not a single tool but a collection of software tools. A data warehouse will collect data from diverse sources into a single database. Using Business Intelligence tools, meaningful insights are drawn from this data. The best thing about "Learn Data Warehousing in 1 Day" is that it is small and can be completed in a day. With this e-book, you will be enough knowledge to contribute and participate in a Data warehouse implementation project. The book covers upcoming and promising technologies like Data Lakes, Data Mart, ELT (Extract Load Transform) amongst others. Following are detailed topics included in the book Table content
Chapter 1: What Is Data Warehouse? What is Data Warehouse? Types of Data

Warehouse Who needs Data warehouse? Why We Need Data Warehouse? Data Warehouse Tools Chapter 2: Data Warehouse Architecture Characteristics of Data warehouse Data Warehouse Architectures Datawarehouse Components Query Tools Chapter 3: ETL Process What is ETL? Why do you need ETL? ETL Process ETL tools Chapter 4: ETL Vs ELT What is ETL? Difference between ETL vs. ELT Chapter 5: Data Modeling What is Data Modelling? Types of Data Models Characteristics of a physical data model Chapter 6: OLAP What is Online Analytical Processing? Types of OLAP systems Advantages and Disadvantages of OLAP Chapter 7: Multidimensional Olap (MOLAP) What is MOLAP? MOLAP Architecture MOLAP Tools Chapter 8: OLAP Vs OLTP What is the meaning of OLAP? What is the meaning of OLTP? Difference between OLTP and OLAP Chapter 9:

Dimensional Modeling What is Dimensional Model? Elements of Dimensional Data Model Attributes Difference between Dimension table vs. Fact table Steps of Dimensional Modelling Rules for Dimensional Modelling Chapter 10: Star and Snowflake Schema What is Multidimensional schemas? What is a Star Schema? What is a Snowflake Schema? Difference between Start Schema and Snowflake Chapter 11: Data Mart What is Data Mart? Type of Data Mart Steps in Implementing a Datamart Chapter 12: Data Mart Vs Data Warehouse What is Data Warehouse? What is Data Mart? Differences between a Data Warehouse and a Data Mart Chapter 13: Data Lake What is Data Lake? Data Lake Architecture Key Data Lake Concepts Maturity stages of Data Lake Chapter 14: Data Lake Vs Data Warehouse What is Data Warehouse? What is Data Lake? Key

Difference between the Data Lake and Data Warehouse Chapter 15: What Is Business Intelligence? What is Business Intelligence Why is BI important? How Business Intelligence systems are implemented? Four types of BI users Chapter 16: Data Mining What is Data Mining? Types of Data Data Mining Process Modelling
The Data Warehouse Toolkit Ralph Kimball 2002 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.
Pentaho Kettle Solutions Matt Casters 2010-09-02 A complete guide to Pentaho Kettle, the Pentaho Data Integration toolset for ETL This practical book is a complete guide to installing, configuring, and managing Pentaho Kettle. If you're a database administrator or developer, you'll first get up

to speed on Kettle basics and how to apply Kettle to create ETL solutions—before progressing to specialized concepts such as clustering, extensibility, and data vault models. Learn how to design and build every phase of an ETL solution. Shows developers and database administrators how to use the open-source Pentaho Kettle for enterprise-level ETL processes (Extracting, Transforming, and Loading data) Assumes no prior knowledge of Kettle or ETL, and brings beginners thoroughly up to speed at their own pace Explains how to get Kettle solutions up and running, then follows the 34 ETL subsystems model, as created by the Kimball Group, to explore the entire ETL lifecycle, including all aspects of data warehousing with Kettle Goes beyond routine tasks to explore how to extend Kettle and scale Kettle solutions using a distributed “cloud” Get the most out of Pentaho Kettle and your

data warehousing with this detailed guide—from simple single table data migration to complex multisystem clustered data integration tasks.

Corporate Information

Factory W. H. Inmon
2002-03-14 The "father of data warehousing" incorporates the latest technologies into his blueprint for integrated decision support systems 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.

Dimensional Modeling: In a Business Intelligence Environment Chuck Ballard 2012-07-31 In this IBM Redbooks publication we describe and demonstrate dimensional data modeling techniques and technology, specifically focused on business intelligence and data warehousing. It is to help the reader understand how to design, maintain, and use a dimensional model for data warehousing that can provide the data access and performance required for business intelligence. Business

intelligence is comprised of a data warehousing infrastructure, and a query, analysis, and reporting environment. Here we focus on the data warehousing infrastructure. But only a specific element of it, the data model - which we consider the base building block of the data warehouse. Or, more precisely, the topic of data modeling and its impact on the business and business applications. The objective is not to provide a treatise on dimensional modeling techniques, but to focus at a more practical level. There is technical content for designing and maintaining such an environment, but also business content. For example, we use case studies to demonstrate how dimensional modeling can impact the business intelligence requirements for your business initiatives. In addition, we provide a detailed discussion on the query aspects of BI and data modeling. For example, we

discuss query optimization and how you can determine performance of the data model prior to implementation. You need a solid base for your data warehousing infrastructure . . . a solid data model.

The Data Warehouse ETL Toolkit Ralph Kimball

2011-04-27 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

The Kimball Group Reader
Ralph Kimball 2010-03-11

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.

A Manager's Guide to Data Warehousing Laura Reeves 2009-05-26 Aimed at helping business and IT managers clearly communicate with each other, this helpful book addresses concerns straight-on and provides practical methods to building a collaborative data warehouse . You'll get clear explanations of the goals and objectives of each stage of the data warehouse

lifecycle while learning the roles that both business managers and technicians play at each stage. Discussions of the most critical decision points for success at each phase of the data warehouse lifecycle help you understand ways in which both business and IT management can make decisions that best meet unified objectives.

Building Event-Driven Microservices

Adam Bellemare 2020-07-02

Organizations today often struggle to balance business requirements with ever-increasing volumes of data. Additionally, the demand for leveraging large-scale, real-time data is growing rapidly among the most competitive digital industries.

Conventional system architectures may not be up to the task. With this practical guide, you'll learn how to leverage large-scale data usage across the business units in your organization using the principles of event-driven

microservices. Author Adam Bellemare takes you through the process of building an event-driven microservice-powered organization. You'll reconsider how data is produced, accessed, and propagated across your organization. Learn powerful yet simple patterns for unlocking the value of this data. Incorporate event-driven design and architectural principles into your own systems. And completely rethink how your organization delivers value by unlocking near-real-time access to data at scale. You'll learn: How to leverage event-driven architectures to deliver exceptional business value The role of microservices in supporting event-driven designs Architectural patterns to ensure success both within and between teams in your organization Application patterns for developing powerful event-driven microservices Components and tooling required to get your microservice

ecosystem off the ground *T-SQL in One Hour a Day, Sams Teach Yourself* Alison Balter 2015-10-16 Master T-SQL database design, development, and administration the easy way—hands-on! In just one hour a day, you'll build all the skills you need to create effective database applications with T-SQL and SQL Server. With this complete tutorial, you'll quickly master the basics and then move on to more advanced features and concepts: Learn the fundamentals of T-SQL from the ground up, one step at a time Succeed with the newest versions of T-SQL, SQL Server, and SQL Server Management Studio Use T-SQL effectively as both an application developer and DBA Master powerful stored procedures, triggers, transactions, and user-defined functions (UDFs) Systematically optimize and secure your SQL Server databases Learn on your own time, at your own pace

No previous T-SQL or database programming experience required Learn how to design efficient, reliable SQL Server databases Define efficient tables, table relationships, fields, and constraints Make the most of T-SQL's SELECT and UPDATE statements Work effectively with simple and complex views and joins Master stored procedure techniques every developer should know Build and use powerful User-Defined Functions (UDFs) Secure databases with authentication, roles, permissions, and principals Configure, maintain, and tune SQL Server for maximum reliability, performance, and value Back up, restore, and audit databases Optimize databases with the SQL Server Profiler, System Monitor, and Index Tuning Wizard Leverage valuable insight and time saving techniques from a world renowned database expert Register your book at

informit.com/register for access to source code, example files, updates, and corrections as they become available.

Database Design for Mere Mortals

Michael James Hernandez 2003 "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* 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 Poolet, 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.

Mastering Data Warehouse Design Claudia Imhoff
2003-08-19 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 approach to data warehousing 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 tough modeling problems, including creating and maintaining keys and modeling calendars, hierarchies, transactions, and data quality

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

Robert Laberge 2011-05-12 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.

Real-Time Analytics Byron Ellis 2014-06-23 Construct a robust end-to-end solution for analyzing and visualizing streaming data Real-time analytics is the hottest topic in data analytics today. In Real-Time Analytics: Techniques to Analyze and Visualize Streaming Data, expert Byron Ellis teaches data analysts technologies to build an effective real-time analytics platform. This platform can then be used

to make sense of the constantly changing data that is beginning to outpace traditional batch-based analysis platforms. The author is among a very few leading experts in the field. He has a prestigious background in research, development, analytics, real-time visualization, and Big Data streaming and is uniquely qualified to help you explore this revolutionary field. Moving from a description of the overall analytic architecture of real-time analytics to using specific tools to obtain targeted results, Real-Time Analytics leverages open source and modern commercial tools to construct robust, efficient systems that can provide real-time analysis in a cost-effective manner. The book includes: A deep discussion of streaming data systems and architectures
Instructions for analyzing, storing, and delivering streaming data
Tips on aggregating data and

working with sets
Information on data warehousing options and techniques
Real-Time Analytics includes in-depth case studies for website analytics, Big Data, visualizing streaming and mobile data, and mining and visualizing operational data flows. The book's "recipe" layout lets readers quickly learn and implement different techniques. All of the code examples presented in the book, along with their related data sets, are available on the companion website.

The Unified Star Schema: An Agile and Resilient Approach to Data Warehouse and Analytics Design Bill Inmon

2020-10-03 Master the most agile and resilient design for building analytics applications: the Unified Star Schema (USS) approach. The USS has many benefits over traditional dimensional modeling. Witness the power of the USS as a single star schema that serves as a

foundation for all present and future business requirements of your organization. Data warehouse legend Bill Inmon and business intelligence innovator, Francesco Puppini, explain step-by-step why the Unified Star Schema is the recommended approach for business intelligence designs today, and show through many examples how to build and use this new solution. This book contains two parts. Part I, Architecture, explains the benefits of data marts and data warehouses, covering how organizations progressed to their current state of analytics, and to the challenges that result from current business intelligence architectures. Chapter 1 covers the drivers behind and the characteristics of the data warehouse and data mart. Chapter 2 introduces dimensional modeling concepts, including fact tables, dimensions, star joins, and

snowflakes. Chapter 3 recalls the evolution of the data mart. Chapter 4 explains Extract, Transform, and Load (ETL), and the value ETL brings to reporting. Chapter 5 explores the Integrated Data Mart Approach, and Chapter 6 explains how to monitor this environment. Chapter 7 describes the different types of metadata within the data warehouse environment. Chapter 8 progresses through the evolution to our current modern data warehouse environment. Part II, the Unified Star Schema, covers the Unified Star Schema (USS) approach and how it solves the challenges introduced in Part I. There are eight chapters within Part II: · Chapter 9, Introduction to the Unified Star Schema: Learn about its architecture and use cases, as well as how the USS approach differs from the traditional approach. · Chapter 10, Loss of Data: Learn about the loss of data and the USS

Bridge. Understand that the USS approach does not create any join, and for this reason, it has no loss of data. · Chapter 11, The Fan Trap: Get introduced to the Oriented Data Model convention, and learn the dangers of a fan trap through an example. Differentiate join and association, and realize that an “in-memory association” is the preferred solution to the fan trap. · Chapter 12, The Chasm Trap: Become familiar with the Cartesian product, and then follow along with an example based on LinkedIn, which illustrates that a chasm trap produces unwanted duplicates. See that the USS Bridge is based on a union, which does not create any duplicates. · Chapter 13, Multi-Fact Queries: Distinguish between multiple facts “with direct connection” versus multiple facts “with no direct connection”. Explore how BI tools are capable of building aggregated virtual rows. ·

Chapter 14, Loops: Learn more about loops and five traditional techniques to solve them. Follow along with an implementation, which will illustrate the solution based on the USS approach. · Chapter 15, Non-Conformed Granularities: Learn about non-conformed granularities, and learn that the Unified Star Schema introduces a solution called “re-normalization”. · Chapter 16, Northwind Case Study. Witness how easy it is to detect the pitfalls of Northwind using the ODM convention. Follow along with an implementation of the USS approach on the Northwind database with various BI tools. [The Data Warehouse Lifecycle Toolkit](#) Ralph Kimball 2011-03-08 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.

Basics of Power BI Modeling

Reza Rad
2020-09-11 I have been dealing with many Power BI challenges in my professional life as a Power BI consultant and a trainer. Challenges normally come as calculation or DAX questions, or sometimes as a performance question. However, after digging deeper into the problem, soon, it will be revealed that the problem is related to a more fundamental challenge; data modeling. If you have a Power BI implementation with many calculation-related or performance-related issues, I strongly suggest looking into your data model because that is where most of the problems start. A good

data model is a great base, which upon that, you can build up many stories of calculations and analysis. A bad data model causes problems on every level that you add upon it, and might sometime cause the whole solution to collapse. Fortunately, data modeling is not rocket science. I explained the basic principles of the data modeling with examples in this book. Use this book as the learning path towards a better data model. Most of the tips mentioned in this book are product-agnostic (such as star-schema, dimension, and fact tables). However, this book is particularly designed and developed for a Power BI product user. This book is for you if you are building a Power BI solution. If your task is only visualizing the existing data, this book might not be needed for you. However, What I have seen in many cases, is that the requirement starts with just visualize the data, and

then more data tables appear, and you get into the tunnel of data modeling without knowing the principles of it. This book is a guide for you through that tunnel.

Star Schema The Complete Reference Christopher Adamson 2010-07-22 The definitive guide to dimensional design for your data warehouse Learn the best practices of dimensional design. Star Schema: The Complete Reference offers in-depth coverage of design principles and their underlying rationales. Organized around design concepts and illustrated with detailed examples, this is a step-by-step guidebook for beginners and a comprehensive resource for experts. This all-inclusive volume begins with dimensional design fundamentals and shows how they fit into diverse data warehouse architectures, including those of W.H. Inmon and

Ralph Kimball. The book progresses through a series of advanced techniques that help you address real-world complexity, maximize performance, and adapt to the requirements of BI and ETL software products. You are furnished with design tasks and deliverables that can be incorporated into any project, regardless of architecture or methodology. Master the fundamentals of star schema design and slow change processing Identify situations that call for multiple stars or cubes Ensure compatibility across subject areas as your data warehouse grows Accommodate repeating attributes, recursive hierarchies, and poor data quality Support conflicting requirements for historic data Handle variation within a business process and correlation of disparate activities Boost performance using derived schemas and aggregates Learn when it's appropriate to adjust

designs for BI and ETL tools
Building the Data Warehouse W. H. Inmon
2002-10-15 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 building data marts, 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 various ways in which 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.

Analyzing Data with Power BI and Power Pivot for Excel

Alberto Ferrari
2017-04-28
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

Kimball's Data

Warehouse Toolkit

Classics Ralph Kimball
2014-02-24 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 structures, providing practical examples of beginning through advanced techniques.