
Access PDF The Data Warehouse ETL Toolkit Practical Techniques For Extracting Cleaning Conforming And Delivering Data

Eventually, you will categorically discover a other experience and talent by spending more cash. yet when? complete you say yes that you require to get those all needs similar to having significantly cash? Why dont you attempt to acquire something basic in the beginning? Thats something that will guide you to understand even more a propos the globe, experience, some places, later than history, amusement, and a lot more?

It is your unconditionally own period to function reviewing habit. in the course of guides you could enjoy now is **The Data Warehouse ETL Toolkit Practical Techniques For Extracting Cleaning Conforming And Delivering Data** below.

8E8GC0 - GWENDOLYN HEIDI

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.

"Ralph's latest book ushers in the second wave of the Internet. . . . Bottom line, this book provides the insight to help companies combine Internet-based business intelligence with the bounty of customer data generated from the internet."--William Schmarzo, Director World Wide Solutions, Sales, and Marketing, IBM NUMA-Q. Receiving over 100 million hits a day, the most popular commercial Websites have an excellent opportunity to collect valuable customer data that can help create better service and improve sales. Companies can use this information to determine buying habits, provide customers with recommendations on new products, and much more. Unfortunately, many companies fail to take full advantage of this deluge of information because they lack the necessary resources to effectively analyze it. In this groundbreaking guide, data warehousing's bestselling author, Ralph Kimball, introduces readers to the Data Webhouse--the marriage of the data warehouse and the Web. If designed and deployed correctly, the Webhouse can become the linchpin of the modern, customer-focused company, providing competitive information essential to managers and strategic decision makers. In this book, Dr. Kimball explains the key elements of the Webhouse and provides detailed guidelines for designing, building, and managing the Webhouse. The results are a business better positioned to stay healthy and competitive. In this book, you'll learn methods for: - Tracking Website user actions - Determining whether a customer is about to switch to a competitor - Determining whether a particular Web ad is working - Capturing data points about customer behavior - Designing the Website to support Webhousing - Building clickstream data marts - Designing the Webhouse user interface - Managing and scaling the Webhouse The companion Website at www.wiley.com/compbooks/kimball provides updates on Webhouse technologies and techniques, as well as links to related

sites and resources.

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.

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 old edition was published in 2002. The current and final edition of this book is *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling*, 3rd Edition which was published in 2013 under ISBN: 9781118530801. The authors begin with fundamental design recommendations and gradually progress step-by-step through increasingly complex scenarios. Clear-cut guidelines for designing dimensional models are illustrated using real-world data warehouse case studies drawn from a variety of business application areas and industries, including: Retail sales and e-commerce Inventory management Procurement Order management Customer relationship management (CRM) Human resources management Accounting Financial services Telecommunications and utilities Education Transportation Health care and insurance By the end of the book, you will have mastered the full range of powerful techniques for designing dimensional databases that are easy to understand and provide fast query response. You will also learn how to create an architected framework that integrates the distributed data warehouse using standardized dimensions and facts.

A thorough update to the industry standard for designing, devel-

oping, 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.

Develop modern solutions with Snowflake's unique architecture and integration capabilities; process bulk and real-time data into a data lake; and leverage time travel, cloning, and data-sharing features to optimize data operations. Key Features: Build and scale modern data solutions using the all-in-one Snowflake platform. Perform advanced cloud analytics for implementing big data and data science solutions. Make quicker and better-informed business decisions by uncovering key insights from your data. Book Description: Snowflake is a unique cloud-based data warehousing platform built from scratch to perform data management on the cloud. This book introduces you to Snowflake's unique architecture, which places it at the forefront of cloud data warehouses. You'll explore the compute model available with Snowflake, and find out how Snowflake allows extensive scaling through the virtual warehouses. You will then learn how to configure a virtual warehouse for optimizing cost and performance. Moving on, you'll get to grips with the data ecosystem and discover how Snowflake integrates with other technologies for staging and loading data. As you progress through the chapters, you will leverage Snowflake's capabilities to process a series of SQL statements using tasks to build data pipelines and find out how you can create modern data solutions and pipelines designed to provide high performance and scalability. You will also get to grips with creating role hierarchies, adding custom roles, and setting default roles for users before covering advanced topics such as data sharing, cloning, and performance optimization. By the end of this Snowflake book, you will be well-versed in Snowflake's architecture for building modern analytical solutions and understand best practices for solving commonly faced problems using practical recipes. What you will learn: Get to grips with data warehousing techniques aligned with Snowflake's cloud architecture. Broaden your skills as a data warehouse designer to cover the Snowflake ecosystem. Transfer skills from on-premise data warehousing to the Snowflake cloud analytics platform. Optimize performance and costs associated with a Snowflake solution. Stage data on object stores and load it into Snowflake. Secure data and share it efficiently for access. Manage transactions and extend Snowflake using stored procedures. Extend cloud data applications using Spark Connector. Who this book is for: This book is for data warehouse developers, data analysts, database administrators, and anyone involved in designing, implementing, and optimizing a Snowflake data warehouse. Knowledge of data warehousing and database and cloud concepts will be useful. Basic familiarity with Snowflake is beneficial, but not necessary.

The Data Vault was invented by Dan Linstedt at the U.S. Depart-

ment 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.

Leverage the power of Microsoft Azure Data Factory v2 to build hybrid data solutions. Key Features: Combine the power of Azure Data Factory v2 and SQL Server Integration Services. Design and enhance performance and scalability of a modern ETL hybrid solution. Interact with the loaded data in data warehouse and data lake using Power BI. Book Description: ETL is one of the essential techniques in data processing. Given data is everywhere, ETL will always be the vital process to handle data from different sources. Hands-On Data Warehousing with Azure Data Factory starts with the basic concepts of data warehousing and ETL process. You will learn how Azure Data Factory and SSIS can be used to understand the key components of an ETL solution. You will go through different services offered by Azure that can be used by ADF and SSIS, such as Azure Data Lake Analytics, Machine Learning and Databrick's Spark with the help of practical examples. You will explore how to design and implement ETL hybrid solutions using different integration services with a step-by-step approach. Once you get to grips with all this, you will use Power BI to interact with data coming from different sources in order to reveal valuable insights. By the end of this book, you will not only learn how to build your own ETL solutions but also address the key challenges that are faced while building them. What you will learn: Understand the key components of an ETL solution using Azure Data Factory and Integration Services. Design the architecture of a modern ETL hybrid solution. Implement ETL solutions for both on-premises and Azure data. Improve the performance and scalability of your ETL solution. Gain thorough knowledge of new capabilities and features added to Azure Data Factory and Integration Services. Who this book is for: This book is for you if you are a software professional who develops and implements ETL solutions using Microsoft SQL Server or Azure cloud. It will be an added advantage if you are a software engineer, DW/ETL architect, or ETL developer, and know how to create a new ETL implementation or enhance an existing one with ADF or SSIS.

DW 2.0: The Architecture for the Next Generation of Data Ware-

housing 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.

As business cycles speed up, many customers gain significant competitive advantage from quicker and more accurate business decision-making by using real data. For many customers, choosing the path to co-locate their transactional and analytical workloads on System z® better leverages their existing investment in hardware, software, and skills. We created a project to address a number of best practice questions on how to manage these newer, analytical type workloads, especially when co-located with traditional transactional workloads. The goal of this IBM® Redbooks® publication is to provide technical guidance and performance trade-offs associated with resource management and potentially DB2® data-sharing in a variety of mixed transactional / data warehouse System z topologies. The term co-location used here and in the rest of the book is specifically defined as the practice of housing both transactional (OLTP) and data warehouse (analytical) workloads within the same System z configuration. We also assumed that key portions of the transactional and data warehouse databases would reside on DB2 for z/OS®. The databases may or may not reside in a DB2 data-sharing environment; we discuss those pros and cons in this book. The intended audience includes DB2 data warehouse architects and practitioners who are facing choices in resource management and system topologies in the data warehouse arena. This specifically includes Business Intelligence (BI) administrators, DB2 database administrators (DBAs) and z/OS performance administrators / systems programmers. In addition, decision makers and architects can utilize this book to assist in making platform and database topology decisions. The book is divided into four parts. Part I, "Introducing the co-location project" covers the System z value proposition and why one should consider System z as the central platform for their data warehousing / business analytics needs. Some topics are risk avoidance via data consolidation, continuous availability, simplified disaster recovery, IBM Smart Analytics Optimizer, reduced network bandwidth requirements, and the unique virtualization and resource management capabilities of System z LPAR, z/VM® and WLM. Part I also provides some of the common System z co-location topologies along with an explanation of the general pros and cons of each. This would be useful input for an architect to understand where a customer is today and where they might consider moving to. Part II, "Project environment" covers the environment, products, workloads, workload drivers, and data models implemented for this study. The environment consisted of

a logically partitioned z10TM 32way, running z/VM, Linux®, and z/OS operating system instances. On those instances we ran products such as z/OS DB2 V9, IBM Cognos® Business Intelligence Version 8.4 for Linux on System z, InfoSphere™ Warehouse for System z, InfoSphere Change Data Capture, z/OS WebSphere® V7, Tivoli® Omegamon for DB2 Performance expert. Utilizing these products we created transactional (OLTP), data warehouse query, and data warehouse refresh workloads. All the workloads were based on an existing web-based transactional Bookstore workload, that's currently utilized for internal testing within the System p® and z labs. While some IBM Cognos BI and ISWz product usage and experiences information is covered in this book, we do not go into the depth typically found in IBM Redbooks publications, since there's another book focused specifically on that 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.

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.

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 Web-

house 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.

Foreword by Mark Stephen LaRow, Vice President of Products, MicroStrategy "A unique and authoritative book that blends recent research developments with industry-level practices for researchers, students, and industry practitioners." Il-Yeol Song, Professor, College of Information Science and Technology, Drexel University ". . . 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.

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.

Geared to IT professionals eager to get into the all-important field of data warehousing, this book explores all topics needed by those who design and implement data warehouses. Readers will learn about planning requirements, architecture, infrastructure, data preparation, information delivery, implementation, and maintenance. They'll also find a wealth of industry examples garnered from the author's 25 years of experience in designing and implementing databases and data warehouse applications for major corporations. Market: IT Professionals, Consultants.

Here is the ideal field guide for data warehousing implementation. This book first teaches you how to build a data warehouse, including defining the architecture, understanding the methodology, gathering the requirements, designing the data models, and creating the databases. Coverage then explains how to populate the data warehouse and explores how to present data to users using reports and multidimensional databases and how to use the

data in the data warehouse for business intelligence, customer relationship management, and other purposes. It also details testing and how to administer data warehouse operation.

Best practices and invaluable advice from world-renowned data warehouse experts In this book, leading data warehouse experts from the Kimball Group share best practices for using the upcoming "Business Intelligence release" of SQL Server, referred to as SQL Server 2008 R2. In this new edition, the authors explain how SQL Server 2008 R2 provides a collection of powerful new tools that extend the power of its BI toolset to Excel and SharePoint users and they show how to use SQL Server to build a successful data warehouse that supports the business intelligence requirements that are common to most organizations. Covering the complete suite of data warehousing and BI tools that are part of SQL Server 2008 R2, as well as Microsoft Office, the authors walk you through a full project lifecycle, including design, development, deployment and maintenance. Features more than 50 percent new and revised material that covers the rich new feature set of the SQL Server 2008 R2 release, as well as the Office 2010 release Includes brand new content that focuses on PowerPivot for Excel and SharePoint, Master Data Services, and discusses updated capabilities of SQL Server Analysis, Integration, and Reporting Services Shares detailed case examples that clearly illustrate how to best apply the techniques described in the book The accompanying Web site contains all code samples as well as the sample database used throughout the case studies The Microsoft Data Warehouse Toolkit, Second Edition provides you with the knowledge of how and when to use BI tools such as Analysis Services and Integration Services to accomplish your most essential data warehousing tasks.

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

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 remain-

ing 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.

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

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 Lawrence Corr 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.

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

Primarily designed as a reference book, simple and effective exercises based upon genuine real-world tasks enable the developer to reduce the time to deliver the results. Presentation of the activities in a recipe format will enable the readers to grasp even the complex concepts with consummate ease. Talend Open Studio Cookbook is principally aimed at relative beginners and intermediate Talend Developers who have used the product to perform some simple integration tasks, possibly via a training course or beginner's tutorials.

Learn the art and science of predictive analytics — techniques that get results Predictive analytics is what translates big data into meaningful, usable business information. Written by a leading expert in the field, this guide examines the science of the underlying algorithms as well as the principles and best practices that govern the art of predictive analytics. It clearly explains the theory behind predictive analytics, teaches the methods, principles, and techniques for conducting predictive analytics projects, and offers tips and tricks that are essential for successful predictive modeling. Hands-on examples and case studies are included. The ability to successfully apply predictive analytics enables businesses to effectively interpret big data; essential for competition today This guide teaches not only the principles of predictive analytics, but also how to apply them to achieve real, pragmatic solutions Explains methods, principles, and techniques for conducting predictive analytics projects from start to finish Illustrates each technique with hands-on examples and includes as series of in-depth case studies that apply predictive analytics to common business scenarios A companion website provides all the data sets used to generate the examples as well as a free trial version of software Applied Predictive Analytics arms data and business analysts and business managers with the tools they need to interpret and capitalize on big data.

This groundbreaking book is the first in the Kimball Toolkit series to be product-specific. Microsoft's BI toolset has undergone significant changes in the SQL Server 2005 development cycle. SQL Server 2005 is the first viable, full-functioned data warehouse

and business intelligence platform to be offered at a price that will make data warehousing and business intelligence available to a broad set of organizations. This book is meant to offer practical techniques to guide those organizations through the myriad of challenges to true success as measured by contribution to business value. Building a data warehousing and business intelligence system is a complex business and engineering effort. While there are significant technical challenges to overcome in successfully deploying a data warehouse, the authors find that the most common reason for data warehouse project failure is insufficient focus on the business users and business problems. In an effort to help people gain success, this book takes the proven Business Dimensional Lifecycle approach first described in best selling *The Data Warehouse Lifecycle Toolkit* and applies it to the Microsoft SQL Server 2005 tool set. Beginning with a thorough description of how to gather business requirements, the book then works through the details of creating the target dimensional model, setting up the data warehouse infrastructure, creating the relational atomic database, creating the analysis services databases, designing and building the standard report set, implementing security, dealing with metadata, managing ongoing maintenance and growing the DW/BI system. All of these steps tie back to the business requirements. Each chapter describes the practical steps in the context of the SQL Server 2005 platform. Intended Audience The target audience for this book is the IT department or service provider (consultant) who is: Planning a small to mid-range data warehouse project; Evaluating or planning to use Microsoft technologies as the primary or exclusive data warehouse server technology; Familiar with the general concepts of data warehousing and business intelligence. The book will be directed primarily at the project leader and the warehouse developers, although everyone involved with a data warehouse project will find the book useful. Some of the book's content will be more technical than the typical project leader will need; other chapters and sections will focus on business issues that are interesting to a database administrator or programmer as guiding information. The book is focused on the mass market, where the volume of data in a single application or data mart is less than 500 GB of raw data. While the book does discuss issues around handling larger warehouses in the Microsoft environment, it is not exclusively, or even primarily, concerned with the unusual challenges of extremely large datasets. About the Authors JOY MUNDY has focused on data warehousing and business intelligence since the early 1990s, specializing in business requirements analysis, dimensional modeling, and business intelligence systems architecture. Joy co-founded InfoDynamics LLC, a data warehouse consulting firm, then joined Microsoft WebTV to develop closed-loop analytic applications and a packaged data warehouse. Before returning to consulting with the Kimball Group in 2004, Joy worked in Microsoft SQL Server product development, managing a team that developed the best practices for building business intelligence systems on the Microsoft platform. Joy began her career as a business analyst in banking and finance. She graduated from Tufts University with a BA in Economics, and from Stanford with an MS in Engineering Economic Systems. WARREN THORNTON has been building data warehousing and business intelligence systems since 1980. Warren worked at Metaphor for eight years, where he managed the consulting organization and implemented many major data warehouse systems. After Metaphor, Warren managed the enterprise-wide data warehouse development at Stanford University. He then co-founded InfoDynamics LLC, a data warehouse consulting firm, with his co-author, Joy Mundy. Warren joined up with WebTV to help build a world class, multi-terabyte customer focused data warehouse before returning to consulting with the Kimball Group. In addition to designing data ware-

houses for a range of industries, Warren speaks at major industry conferences and for leading vendors, and is a long-time instructor for Kimball University. Warren holds an MBA in Decision Sciences from the University of Pennsylvania's Wharton School, and a BA in Communications Studies from the University of Michigan. RALPH KIMBALL, PH.D., has been a leading visionary in the data warehouse industry since 1982 and is one of today's most internationally well-known authors, speakers, consultants, and teachers on data warehousing. He writes the "Data Warehouse Architect" column for *Intelligent Enterprise* (formerly *DBMS*) magazine.

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.

Market_Desc: · Data warehouse Designers· Data warehouse Architects· Data warehouse Developers· Data warehouse Managers
Special Features: · The current first edition has sold more than 72,000 copies, generating net revenue of more than \$2.5 million· The methods described in this book have been adopted by almost all leading data warehouse vendors· Ralph Kimball and his co-authors are recognized as the driving thought leaders in the data warehousing industry; there is no direct competition· The authors actively promote this methodology in training and consulting worldwide and in their writing in magazines and online
About The Book: The book covers best practices from data warehouse project inception through on-going program management. About 30 to 40% of the content in the book is updated and new. This revised tutorial covers major lifecycle topics such as dimensional modeling, tech architecture, ETL, BI etc. It is targeted at both novice and experienced data warehouse professionals.

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, spa-

tial, 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.

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 Ware-

house 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

This book is written in a recipe-based style packed full of practical tips and techniques to help you analyse multidimensional data stored in SSAS 2012 cubes. If you need to master MDX queries in SSAS, then this book is for you! If you are a Microsoft SQL Server Analysis Services developer and want to improve your solutions using MDX, then this book is for you. This book is also an essential resource for report developers who need to access the multidimensional cubes through the MDX language. The book assumes you have some basic working knowledge of MDX and a basic understanding of dimensional modelling and cube design.

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

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)

This book presents the first comparative review of the state of the art and the best current practices of data warehouses. It covers source and data integration, multidimensional aggregation, query optimization, metadata management, quality assessment, and design optimization. A conceptual framework is presented by which the architecture and quality of a data warehouse can be assessed and improved using enriched metadata management combined with advanced techniques from databases, business modeling, and artificial intelligence.

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 sub-

ject 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