

Agile Data Warehouse Design Collaborative Dimensional Modeling From Whiteboard To Star Schema

Data Modeling Essentials, Third Edition, covers the basics of data modeling while focusing on developing a facility in techniques, rather than a simple familiarization with "the rules". In order to enable students to apply the basics of data modeling to real models, the book addresses the realities of developing systems in real-world situations by assessing the merits of a variety of possible solutions as well as using language and diagramming methods that represent industry practice. This revised edition has been given significantly expanded coverage and reorganized for greater reader comprehension even as it retains its distinctive hallmarks of readability and usefulness. Beginning with the basics, the book provides a thorough grounding in theory before guiding the reader through the various stages of applied data modeling and database design. Later chapters address advanced subjects, including business rules, data warehousing, enterprise-wide modeling and data management. It includes an entirely new section discussing the development of logical and physical modeling, along with new material describing a powerful technique for model verification. It also provides an excellent resource for additional lectures and exercises. This text is the ideal reference for data modelers, data architects, database designers, DBAs, and systems analysts, as well as undergraduate and graduate-level students looking for a real-world perspective. Thorough coverage of the fundamentals and relevant theory. Recognition and support for the creative side of the process. Expanded coverage of applied data modeling includes new chapters on logical and physical database design. New material describing a powerful technique for model verification. Unique coverage of the practical and human aspects of modeling, such as working with business specialists, managing change, and resolving conflict.

An exciting new approach to dimensional data warehouse design, invented by the author This how-to guide shows you how a revolutionary new approach to dimensional modeling invented by the author can be used throughout the data warehouse analysis and design phases to improve communication between IT and business users-and consequently-produce better results. Featuring proven techniques devised by the author, this book introduces Business Event Analysis and Modeling (BEAM), a powerful, yet simple, technique for modeling directly with business users. You'll learn how to run facilitated BEAM sessions with users to create consensus on data definitions and a sense of business ownership of the data warehouse design. Demonstrates how to use Business Event Analysis and Modeling (BEAM) to readily capture the 5W1H combinations that represent the unique business events of any organization Presents common reusable modeling techniques for handling the complexities of people and groups (the "who"), products and services (the "what"), geography (the "where") for each dimension type Offers extensive, detailed examples for translating each BEAM model into a dimensional (star schema) model The companion website will include all sample models, document templates, and SQL code examples These proven techniques will help you accomplish your end goal: ensuring that the data warehouse design meet the real needs of users.

World-renowned economist Klaus Schwab, Founder and Executive Chairman of the World Economic Forum, explains that we have an opportunity to shape the fourth industrial revolution, which will fundamentally alter how we live and work. Schwab argues that this revolution is different in scale, scope and complexity from any that have come before. Characterized by a range of new technologies that are fusing the physical, digital and biological worlds, the developments are affecting all disciplines, economies, industries and governments, and even challenging ideas about what it means to be human. Artificial intelligence is already all around us, from supercomputers, drones and virtual assistants to 3D printing, DNA sequencing, smart thermostats, wearable sensors and microchips smaller than a grain of sand. But this is just the beginning: nanomaterials 200 times stronger than steel and a million times thinner than a strand of hair and the first transplant of a 3D printed liver are already in development. Imagine "smart factories" in which global systems of manufacturing are coordinated virtually, or implantable mobile phones made of biosynthetic materials. The fourth industrial revolution, says Schwab, is more significant, and its ramifications more profound, than in any prior period of human history. He outlines the key technologies driving this revolution and discusses the major impacts expected on government, business, civil society and individuals. Schwab also offers bold ideas on how to harness these changes and shape a better future--one in which technology empowers people rather than replaces them; progress serves society rather than disrupts it; and in which innovators respect moral and ethical boundaries rather than cross them. We all have the opportunity to contribute to developing new frameworks that advance progress.

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.

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 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.

Work with petabyte-scale datasets while building a collaborative, agile workplace in the process. This practical book is the canonical reference to Google BigQuery, the query engine that lets you conduct interactive analysis of large datasets. BigQuery enables enterprises to efficiently store, query, ingest, and learn from their data in a convenient framework. With this book, you'll examine how to analyze data at scale to derive insights from large datasets efficiently. Valliappa Lakshmanan, tech lead for Google Cloud Platform, and Jordan Tigani, engineering director for the BigQuery team, provide best practices for modern data warehousing within an autoscaled, serverless public cloud. Whether you want to explore parts of BigQuery you're not familiar with or prefer to focus on specific tasks, this reference is indispensable.

Written by a team of global experts, this book explains how to design next-generation data warehouses using a structured approach inspired by the modern principles of software engineering. All of the design techniques featured in the book are placed within a practical framework so that solutions can be applied immediately. Data Warehouse Design: Modern Principles and Methodologies explains how to transform databases into data warehouses, allowing organizations to integrate and analyze data across the enterprise in order to make informed business decisions. Readers can access a CASE (computer aided software engineering) tool developed by the authors from www.qbx-tool.com which implements the discussed methodologies in a practical way.

Building upon his earlier book that detailed agile data warehousing programming techniques for the Scrum master, Ralph's latest work illustrates the agile interpretations of the remaining software engineering disciplines: Requirements management benefits from streamlined templates that not only define projects quickly, but ensure nothing essential is overlooked. Data engineering receives two new "hyper modeling" techniques, yielding data warehouses that can be easily adapted when requirements change without having to invest in ruinously expensive data-conversion programs. Quality assurance advances with not only a stereoscopic top-down and bottom-up planning method, but also the incorporation of the latest in automated test engines. Use this step-by-step guide to deepen your own application development skills through self-study, show your teammates the world's fastest and most reliable techniques for creating business intelligence systems, or ensure that the IT department working for you is building your next decision support system the right way. Learn how to quickly define scope and architecture before programming starts Includes techniques of process and data engineering that enable iterative and incremental delivery Demonstrates how to plan and execute quality assurance plans and includes a guide to continuous integration and automated regression testing Presents program management strategies for coordinating multiple agile data mart projects so that over time an enterprise data warehouse emerges Use the provided 120-day road map to establish a robust, agile data warehousing program

The world of data warehousing is changing. Big Data & Agile are hot topics. But companies still need to collect, report, and analyze their data. Usually this requires some form of data warehousing or business intelligence system. So how do we do that in the modern IT landscape in a way that allows us to be agile and either deal directly or indirectly with unstructured and semi structured data?The Data Vault System of Business Intelligence provides a method and approach to modeling your enterprise data warehouse (EDW) that is agile, flexible, and scalable. This book will give you a short introduction to Agile Data Engineering for Data Warehousing and Data Vault 2.0. I will explain why you should be trying to become Agile, some of the history and rationale for Data Vault 2.0, and then show you the basics for how to build a data warehouse model using the Data Vault 2.0 standards. In addition, I will cover some details about the Business Data Vault (what it is) and then how to build a virtual Information Mart off your Data Vault and Business Vault using the Data Vault 2.0 architecture. So if you want to start learning about Agile Data Engineering with Data Vault 2.0, this book is for you.

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

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.

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

"Each chapter is... a practice run for the way we all ought to design our data marts and hence our data warehouses."-Ralph Kimball, from the Foreword. Let the experts show you how to customize data warehouse designs for real business needs in Data Warehouse Design Solutions. To effectively design a data warehouse, you have to understand its many business uses. This guidebook shows you how business managers in different corporate functions actually use data warehouses to make decisions. You'll get a rich set of data warehouse designs that flow from realistic business cases. Two top experts show you how to customize your data warehouse designs for real-life business needs including: * Sales and marketing * Production and inventory management * Budgeting and financial reporting * Quality control * Product delivery and fulfillment * Strategic business analysis such as determining market share, rates of return on investment, and other key analytic ratios. CD-ROM includes All sample data warehouse designs with accompanying preformatted reports in HTML for specific business uses such as marketing, sales, and financial analysis.

The data vault methodology presents a unique opportunity to model the enterprise data warehouse using the same automation principles applicable in today's software delivery, continuous integration, continuous delivery and continuous deployment while still maintaining the standards expected for governing a corporation's most valuable asset: data. This book provides at first the landscape of a modern architecture and then as a thorough guide on how to deliver a data model that flexes as the enterprise flexes, the data vault. Whether the data is structured, semi-structured or even unstructured one thing is clear, there is always a model either applied early (schema-on-write) or applied late (schema-on-read). Today's focus on data governance requires that we know what we retain about our customers, the data vault provides that focus by delivering a methodology focused on all aspects about the customer and provides some of the best practices for modern day data compliance. The book will delve into every data vault modelling artefact, its automation with sample code, raw vault, business vault, testing framework, a build framework, sample data vault models, how to build automation patterns on top of a data vault and even offer an extension of data vault that provides automated timeline correction, not to mention variation of data vault designed to provide audit trails, metadata control and integration with agile delivery tools.

The first book to cover Agile Modeling, a new modeling technique created specifically for XP projects eXtreme Programming (XP) has created a buzz in the software development community- much like Design Patterns did several years ago. Although XP presents a methodology for faster software development, many developers find that XP does not allow for modeling time, which is critical to ensure that a project meets its proposed requirements. They have also found that standard modeling techniques that use the Unified Modeling Language (UML) often do not work with this methodology. In this innovative book, Software Development columnist Scott Ambler presents Agile Modeling (AM)-a technique that he created for modeling XP projects using pieces of the UML and Rational's Unified Process (RUP). Ambler clearly explains AM, and shows readers how to incorporate AM, UML, and RUP into their development projects with the help of numerous case studies integrated throughout the book. AM was created by the author for modeling XP projects- an element lacking in the original XP design The XP community and its creator have embraced AM, which should give this book strong market acceptance Companion Web site at www.agilemodeling.com features updates, links to XP and AM resources, and ongoing case studies about agile modeling.

The final edition of the incomparable data warehousing and business intelligence reference, updated and expanded The Kimball Group Reader, Remastered Collection is the essential reference for data warehouse and business intelligence design, packed with best practices, design tips, and valuable insight from industry pioneer Ralph Kimball and the Kimball Group. This Remastered Collection represents decades of expert advice and mentoring in data warehousing and business intelligence, and is the final work to be published by the Kimball Group. Organized for quick navigation and easy reference, this book contains nearly 20 years of experience on more than 300 topics, all fully up-to-date and expanded with 65 new articles. The

discussion covers the complete data warehouse/business intelligence lifecycle, including project planning, requirements gathering, system architecture, dimensional modeling, ETL, and business intelligence analytics, with each group of articles prefaced by original commentaries explaining their role in the overall Kimball Group methodology. Data warehousing/business intelligence industry's current multi-billion dollar value is due in no small part to the contributions of Ralph Kimball and the Kimball Group. Their publications are the standards on which the industry is built, and nearly all data warehouse hardware and software vendors have adopted their methods in one form or another. This book is a compendium of Kimball Group expertise, and an essential reference for anyone in the field. Learn data warehousing and business intelligence from the field's pioneers Get up to date on best practices and essential design tips Gain valuable knowledge on every stage of the project lifecycle Dig into the Kimball Group methodology with hands-on guidance Ralph Kimball and the Kimball Group have continued to refine their methods and techniques based on thousands of hours of consulting and training. This Remastered Collection of The Kimball Group Reader represents their final body of knowledge, and is nothing less than a vital reference for anyone involved in the field.

You want the rigor of good data architecture at the speed of agile? Then this is the missing link - your step-by-step guide to Data Vault success. Success with a Data Vault starts with the business and ends with the business. Sure, there's some technical stuff in the middle, and it is absolutely essential - but it's not sufficient on its own. This book will help you shape the business perspective, and weave it into the more technical aspects of Data Vault modeling. You can read the foundational books and go on courses, but one massive risk still remains. Dan Linstedt, the founder of the Data Vault, very clearly directs those building a Data Vault to base its design on an "enterprise ontology". And Hans Hultgren similarly stresses the importance of the business concepts model. So it's important. We get that. But: What on earth is an enterprise ontology/business concept model, 'cause I won't know if I've got one if I don't know what I'm looking for? If I can't find one, how do I get my hands on such a thing? Even if I have one of these wonderful things, how do I apply it to get the sort of Data Vault that's recommended? It's actually not as hard as some would fear to answer all of these questions, and it's certainly worth the effort. This book just might save you a world of pain. It's a supplement to other material on Data Vault modeling, but it's the vital missing link to finding simplicity for Data Vault success.

Explore the modern market of data analytics platforms and the benefits of using Snowflake computing, the data warehouse built for the cloud. With the rise of cloud technologies, organizations prefer to deploy their analytics using cloud providers such as Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform. Cloud vendors are offering modern data platforms for building cloud analytics solutions to collect data and consolidate into single storage solutions that provide insights for business users. The core of any analytics framework is the data warehouse, and previously customers did not have many choices of platform to use. Snowflake was built specifically for the cloud and it is a true game changer for the analytics market. This book will help onboard you to Snowflake, present best practices to deploy, and use the Snowflake data warehouse. In addition, it covers modern analytics architecture and use cases. It provides use cases of integration with leading analytics software such as Matillion ETL, Tableau, and Databricks. Finally, it covers migration scenarios for on-premise legacy data warehouses. What You Will Learn Know the key functionalities of Snowflake Set up security and access with cluster Bulk load data into Snowflake using the COPY command Migrate from a legacy data warehouse to Snowflake integrate the Snowflake data platform with modern business intelligence (BI) and data integration tools Who This Book Is For Those working with data warehouse and business intelligence (BI) technologies, and existing and potential Snowflake users

This publication assesses progress towards Sustainable Development Goal 4 (SDG 4) on education and its ten targets, as well as other related education targets in the SDG agenda. It addresses inclusion in education, drawing attention to all those excluded from education, because of background or ability. The report is motivated by the explicit reference to inclusion in the 2015 Incheon Declaration, and the call to ensure an inclusive and equitable quality education in the formulation of SDG 4, the global goal for education. It reminds us that, no matter what argument may be built to the contrary, we have a moral imperative to ensure every child has a right to an appropriate education of high quality.

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

Agile Data Warehouse Design Collaborative Dimensional Modeling, from Whiteboard to Star Schema DecisionOne Consulting

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

Azure Data Lake Gen 2 to feed your data warehouse Who This Book Is For Data warehouse or ETL/ELT developers who wish to implement a data warehouse project in the Azure cloud, and developers currently working in on-premise environments who want to move to the cloud, and for developers with Azure experience looking to tighten up their implementation and consolidate their knowledge

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

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.

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 innaccurate data, transforming the remaining data into correctly formatted data structures, and then physically loading them into the data warehouse. This book provides complete coverage of proven, time-saving ETL techniques. It begins with a quick overview of ETL fundamentals and the role of the ETL development team. It then quickly moves into an overview of the ETL data structures, both relational and dimensional. The authors show how to build useful dimensional stuctures, providing practical examples of beginning through advanced techniques.

Using Agile methods, you can bring far greater innovation, value, and quality to any data warehousing (DW), business intelligence (BI), or analytics project. However, conventional Agile methods must be carefully adapted to address the unique characteristics of DW/BI projects. In Agile Analytics, Agile pioneer Ken Collier shows how to do just that. Collier introduces platform-agnostic Agile solutions for integrating infrastructures consisting of diverse operational, legacy, and specialty systems that mix commercial and custom code. Using working examples, he shows how to manage analytics development teams with widely diverse skill sets and how to support enormous and fast-growing data volumes. Collier's techniques offer optimal value whether your projects involve "back-end" data management, "front-end" business analysis, or both. Part I focuses on Agile project management techniques and delivery team coordination, introducing core practices that shape the way your Agile DW/BI project community can collaborate toward success Part II presents technical methods for enabling continuous delivery of business value at production-quality levels, including evolving superior designs; test-driven DW development; version control; and project automation Collier brings together proven solutions you can apply right now--whether you're an IT decision-maker, data warehouse professional, database administrator, business intelligence specialist, or database developer. With his help, you can mitigate project risk, improve business alignment, achieve better results--and have fun along the way.

Data science teams looking to turn research into useful analytics applications require not only the right tools, but also the right approach if they're to succeed. With the revised second edition of this hands-on guide, up-and-coming data scientists will learn how to use the Agile Data Science development methodology to build data applications with Python, Apache Spark, Kafka, and other tools. Author Russell Journey demonstrates how to compose a data platform for building, deploying, and refining analytics applications with Apache Kafka, MongoDB, Elasticsearch, d3.js, scikit-learn, and Apache Airflow. You'll learn an iterative approach that lets you quickly change the kind of analysis you're doing, depending on what the data is telling you. Publish data science work as a web application, and affect meaningful change in your organization. Build value from your data in a series of agile sprints, using the data-value pyramid Extract features for statistical models from a single dataset Visualize data with charts, and expose different aspects through interactive reports Use historical data to predict the future via classification and

regression Translate predictions into actions Get feedback from users after each sprint to keep your project on track

Contains a six-stage plan for starting new warehouse projects and guiding programmers step-by-step until they become a world-class, Agile development team. It describes also how to avoid or contain the fierce opposition that radically new methods can encounter from the traditionally-minded IS departments found in many large companies.

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.

As data management and integration continue to evolve rapidly, storing all your data in one place, such as a data warehouse, is no longer scalable. In the very near future, data will need to be distributed and available for several technological solutions. With this practical book, you'll learn how to migrate your enterprise from a complex and tightly coupled data landscape to a more flexible architecture ready for the modern world of data consumption. Executives, data architects, analytics teams, and compliance and governance staff will learn how to build a modern scalable data landscape using the Scaled Architecture, which you can introduce incrementally without a large upfront investment. Author Pietheine Strengholt provides blueprints, principles, observations, best practices, and patterns to get you up to speed. Examine data management trends, including technological developments, regulatory requirements, and privacy concerns Go deep into the Scaled Architecture and learn how the pieces fit together Explore data governance and data security, master data management, self-service data marketplaces, and the importance of metadata

This book helps you to understand Snowflake's unique architecture and ecosystem that places it at the forefront of cloud data warehouses. The recipes present in this book will enable you to develop proficiency in managing data on Snowflake and learn Snowflake's novel features such as data sharing, cloning, and time travel.

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

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.

Data Modeling for Agile Data Warehouse using Data Vault Modeling Approach. Includes Enterprise Data Warehouse Architecture. This is a complete guide to the data vault data modeling approach. The book also includes business and program considerations for the agile data warehousing and business intelligence program. There are over 200 diagrams and figures concerning modeling, core business concepts, architecture, business alignment, semantics, and modeling comparisons with 3NF and Dimensional modeling.

You have to make sense of enormous amounts of data, and while the notion of "agile data warehousing might sound tricky, it can yield as much as a 3-to-1 speed advantage while cutting project costs in half. Bring this highly effective technique to your organization with the wisdom of agile data warehousing expert Ralph Hughes. Agile Data Warehousing Project Management will give you a thorough introduction to the method as you would practice it in the project room to build a serious "data mart. Regardless of where you are today, this step-by-step implementation guide will prepare you to join or even lead a team in visualizing, building, and validating a single component to an enterprise data warehouse. Provides a thorough grounding on the mechanics of Scrum as well as practical advice on keeping your team on track Includes strategies for getting accurate and actionable requirements from a team's business partner Revolutionary estimating techniques that make forecasting labor far more understandable and accurate Demonstrates a blends of Agile methods to simplify team management and synchronize inputs across IT specialties Enables you and your teams to start simple and progress steadily to world-class performance levels

This is the digital version of the printed book (Copyright © 1996). Learning the basics of a modeling technique is not the same as learning how to use and apply it. To develop a data model of an organization is to gain insights into its nature that do not come easily. Indeed, analysts are often expected to understand subtleties of an organization's structure that may have evaded people who have worked there for years. Here's help for those analysts who have learned the basics of data modeling (or "entity/relationship modeling") but who need to obtain the insights required to prepare a good model of a real business. Structures common to many types of business are analyzed in areas such as accounting, material requirements planning, process manufacturing, contracts, laboratories, and documents. In each chapter, high-level data models are drawn from the following business areas: The Enterprise and Its World The Things of the Enterprise Procedures and Activities Contracts Accounting The Laboratory Material Requirements Planning Process Manufacturing Documents Lower-Level Conventions

Hundreds of organizations around the world have already benefited from Disciplined Agile Delivery (DAD). Disciplined Agile (DA) is the only comprehensive tool kit available for guidance on building high-performance agile teams and optimizing your way of working (WoW). As a hybrid of all the leading agile and lean approaches, it provides hundreds of strategies to help you make better decisions within your agile teams, balancing self-organization with the realities and constraints of your unique enterprise context. The highlights of this handbook include:

- As the official source of knowledge on DAD, it includes greatly improved and enhanced strategies with a revised set of goal diagrams based upon learnings from applying DAD in the field.
- It is an essential handbook to help coaches and teams make better decisions in their daily work, providing a wealth of ideas for experimenting with agile and lean techniques while providing specific guidance and trade-offs for those "it depends" questions.
- It makes a perfect study guide for Disciplined Agile certification. Why "fail fast" (as our industry likes to recommend) when you can learn quickly on your journey to high performance? With this handbook, you can make better decisions based upon proven, context-based strategies, leading to earlier success and better outcomes.

Mining big data requires a deep investment in people and time. How can you be sure you're building the right models? With this hands-on book, you'll learn a flexible toolset and methodology for building effective analytics applications with Hadoop. Using lightweight tools such as Python, Apache Pig, and the D3.js library, your team will create an agile environment for exploring data, starting with an example application to mine your own email inboxes. You'll learn an iterative approach that enables you to quickly change the kind of analysis you're doing, depending on what the data is telling you. All example code in this book is available as working Heroku apps. Create analytics applications by using the agile big data development methodology Build value from your data in a series of agile sprints, using the data-value stack Gain insight by using several data structures to extract multiple features from a single dataset Visualize data with charts, and expose different aspects through interactive reports Use historical data to predict the future, and translate predictions into action Get feedback from users after each sprint to keep your project on track

Describes Agile Modeling Driven Design (AMDD) and Test-Driven Design (TDD) approaches, database refactoring, database encapsulation strategies, and tools that support evolutionary techniques Agile software developers often use object and relational database (RDB) technology together and as a result must overcome the impedance mismatch The author covers techniques for mapping objects to RDBs and for implementing concurrency control, referential integrity, shared business logic, security access control, reports, and XML An agile foundation describes fundamental skills that all agile software developers require, particularly Agile DBAs Includes object modeling, UML data modeling, data normalization, class normalization, and how to deal with legacy databases Scott W. Ambler is author of Agile Modeling (0471202827), a contributing editor with Software Development (www.sdmagazine.com), and a featured speaker at software conferences worldwide

[Copyright: 8cb05b9cdeb92bfb8103f99ed528686a](#)