

Agile Data Warehousing For The Enterprise A Guide For Solution Architects And Project Leaders

Do You Know If Your Data Warehouse Flexible, Scalable, Secure and Will It Stand The Test Of Time And Avoid Being Part Of The Dreaded "Life Cycle"? The Data Vault took the Data Warehouse world by storm when it was released in 2001. Some of the world's largest and most complex data warehouse situations understood the value it gave especially with the capabilities of unlimited scaling, flexibility and security. Here is what industry leaders say about the Data Vault "The Data Vault is the optimal choice for modeling the EDW in the DW 2.0 framework" - Bill Inmon, The Father of Data Warehousing "The Data Vault is foundationally strong and an exceptionally scalable architecture" - Stephen Brobst, CTO, Teradata "The Data Vault should be considered as a potential standard for RDBMS-based analytic data management by organizations looking to achieve a high degree of flexibility, performance and openness" - Doug Lane, Deloitte Analytics Institute "I applaud Dan's contribution to the body of Business Intelligence and Data Warehousing knowledge and recommend this book be read by both data professionals and end users" - Howard Dresner, From the Foreword - Speaker, Author, Leading Research Analyst and Advisor You have in your hands the work, experience and testing of 2 decades of building data warehouses. The Data Vault model and methodology has proven itself in hundreds (perhaps thousands) of solutions in Insurance, Crime-Fighting, Defense, Retail, Finance, Banking, Power, Energy, Education, High-Tech and many more. Learn the techniques and implement them and learn how to build your Data Warehouse faster than you have ever done before while designing it to grow and scale no matter what you throw at it. Ready to "Super Charge Your Data Warehouse"? 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. Refactoring has proven its value in a wide range of development projects—helping software professionals improve system designs, maintainability, and performance. Now, for the first time, leading agile methodologist Scott Ambler and renowned consultant Pramodkumar Sadalage introduce powerful refactoring techniques specifically designed for database systems. Ambler and Sadalage demonstrate how small changes to table structures, data, stored procedures, and triggers can significantly enhance virtually any database design—without changing semantics. You'll learn how to evolve database schemas in step with source code—and become far more effective in projects relying on iterative, agile methodologies. This comprehensive guide and reference helps you overcome the practical obstacles to refactoring real-world databases by covering every fundamental concept underlying database refactoring. Using start-to-finish examples, the authors walk you through refactoring simple standalone database applications as well as sophisticated multi-application scenarios. You'll master every task involved in refactoring database schemas, and discover best practices for deploying refactorings in even the most complex production environments. The second half of this book systematically covers five major categories of database refactorings. You'll learn how to use refactoring to enhance database structure, data quality, and referential integrity; and how to refactor both architectures and methods. This book provides an extensive set of examples built with Oracle and Java and easily adaptable for other languages, such as C#, C++, or VB.NET, and other databases, such as DB2, SQL Server, MySQL, and Sybase. Using this book's techniques and examples, you can reduce waste, rework, risk, and cost—and build database systems capable of evolving smoothly, far into the future. 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 Olschmike 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 Agile Analytics Agile Data Warehousing Project Management NoSQL Distilled A Value-driven Approach to Business Intelligence and Data Warehousing A Guide for Solution Architects and Project Leaders 3 Volume Set Inmon Versus Kimball

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

Praise for Successful Business Intelligence "If you want to be an analytical competitor, you've got to go well beyond business intelligence technology. Cindi Howson has wrapped up the needed advice on technology, organization, strategy, and even culture in a neat package. It's required reading for quantitatively oriented strategists and the technologists who support them." --Thomas H. Davenport, President's Distinguished Professor, Babson College and co-author, Competing on Analytics "When used strategically, business intelligence can help companies transform their organization to be more agile, more competitive, and more profitable. Successful Business Intelligence offers valuable guidance for companies looking to embark upon their first BI project as well as those hoping to maximize their current deployments." --John Schwarz, CEO, Business Objects "A thoughtful, clearly written, and carefully researched examination of all facets of business intelligence that your organization needs to know to run its business more intelligently and exploit information to its fullest extent." --Wayne Eckerson, Director, TDWI Research "Using real-world examples, Cindi Howson shows you how to use business intelligence to improve the performance, and the quality, of your company." --Bill Baker, Distinguished Engineer & GM, Business Intelligence Applications, Microsoft Corporation "This book outlines the key steps to make BI an integral part of your company's culture and demonstrates how your company can use BI as a competitive differentiator." --Robert VanHees, CFO, Corporate Express "Given the trend to expand the business analytics user base, organizations are faced with a number of challenges that affect the success rate of these projects. This insightful book provides practical advice on improving that success rate." --Dan Vesset, Vice President, Business Analytics Solution Research, IDC

Between the high-level concepts of business intelligence and the nitty-gritty instructions for using vendors' tools lies the essential, yet poorly-understood layer of architecture, design and process. Without this knowledge, Big Data is belittled - projects flounder, are late and go over budget. Business Intelligence Guidebook: From Data Integration to Analytics shines a bright light on an often neglected topic, arming you with the knowledge you need to design rock-solid business intelligence and data integration processes. Practicing consultant and adjunct BI professor Rick Sherman takes the guesswork out of creating systems that are cost-effective, reusable and essential for transforming raw data into valuable information for business decision-makers. After reading this book, you will be able to design the overall architecture for functioning business intelligence systems with the supporting data warehousing and data-integration applications. You will have the information you need to get a project launched, developed, managed and delivered on time and on budget - turning the deluge of data into actionable information that fuels business knowledge. Finally, you'll give your career a boost by demonstrating an essential knowledge that puts corporate BI projects on a fast-track to success. Provides practical guidelines for building successful BI, DW and data integration solutions. Explains underlying BI, DW and data integration design, architecture and processes in clear, accessible language. Includes the complete project development lifecycle that can be applied at large enterprises as well as at small to medium-sized businesses Describes best practices and pragmatic approaches so readers can put them into action. Companion website includes templates and examples, further discussion of key topics, instructor materials, and references to trusted industry sources.

Hundreds of organizations around the world have already benefited from DisciplinedAgile 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.

Relentlessly Practical Tools for Data Warehousing and Business Intelligence

The Data Warehouse Toolkit

Invaluable Data Modeling Rules to Implement Your Data Vault

Data Warehousing Fundamentals

Building the Data Warehouse

An Introduction to Agile Data Engineering Using Data Vault 2. 0

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

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

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

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.

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.

Agile Data Warehousing for the Enterprise

Mastering Data Warehouse Aggregates

A Small Agile Team's Journey from Scrum to Disciplined DevOps

From Data Integration to Analytics

Lessons Learned for the Data-Driven Business

Introduction to Disciplined Agile Delivery 2nd Edition

Extreme Scoping

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. 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

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.

PLEASE PROVIDE COURSE INFORMATION PLEASE PROVIDE

Applied Data Science

Data Warehousing in the Age of Big Data

Refactoring Databases

Successful Business Intelligence: Secrets to Making BI a Killer App

Refactoring SQL Applications

Agile – Data Warehousing: the Backbone of Business Intelligence

Mastering Data Warehouse Design

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.

Data Warehousing in the Age of the Big Data will help you and your organization make the most of unstructured data with your existing data warehouse. As Big Data continues to revolutionize how we use data, it doesn't have to create more confusion. Expert author Krish Krishnan helps you make sense of how Big Data fits into the world of data warehousing in clear and concise detail. The book is presented in three distinct parts. Part 1 discusses Big Data, its technologies and use cases from early adopters. Part 2 addresses data warehousing, its shortcomings, and new architecture options, workloads, and integration techniques for Big Data and the data warehouse. Part 3 deals with data governance, data visualization, information life-cycle management, data scientists, and implementing a Big Data-ready data warehouse. Extensive appendixes include case studies from vendor implementations and a special segment on how we can build a healthcare information factory. Ultimately, this book will help you navigate through the complex layers of Big Data and data warehousing while providing you information on how to effectively think about using all these technologies and the architectures to design the next-generation data warehouse. Learn how to leverage Big Data by effectively integrating it into your data warehouse. Includes real-world examples and use cases that clearly demonstrate Hadoop, NoSQL, HBASE, Hive, and other Big Data technologies Understand how to optimize and tune your current data warehouse infrastructure and integrate newer infrastructure matching data processing workloads and requirements 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. 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.

Solutions for Star Schema Performance

Data Warehousing

Better Data Modeling

Modeling the Agile Data Warehouse with Data Vault

Super Charge Your Data Warehouse

Relational and Dimensional Techniques

An Agile Approach to Enterprise Data Warehousing and Business Intelligence

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.

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

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.

"This all day workshop will provide you the background and principles to use agile in your data warehouse and business intelligence projects. It will introduce you to a modern method for agile data modeling, Data Vault 2.0, and provide you with a detailed, real world case study. At the end we will talk about how the cloud has changed everything and how you can enable your agile data warehouse by using a modern data warehouse as a service (DWaaS) built in the cloud, for the cloud."--Resource description page.

A Thesis in Computer Science

Volume 1

Incorporating Agile Principles

Trends and Advances in Information Systems and Technologies

Kimball's Data Warehouse Toolkit Classics

The Kimball Group Reader

Building a Scalable Data Warehouse with Data Vault 2.0

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 need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational "NoSQL" databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage's pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Polyglot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access.

What is agile data warehousing? -- Iterative development in a nutshell -- Streamlining project management -- Authoring better user stories -- Deriving initial project backlogs -- Developer stories for data integration -- Estimating and segmenting projects -- Adapting agile for data warehousing -- Starting and scaling agile data warehousing.

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

Business Intelligence Guidebook

Choose your Wow

The Complete Guide to Dimensional Modeling

Agile Data Warehousing in the Cloud

The Data Model Resource CD

Business Intelligence Systems Using Scrum

Agile Data Warehouse Design

Do your business intelligence (BI) projects take too long to deliver? Is the value of the deliverables less than satisfactory? Do these projects propagate poor data management practices? If you screamed [yes] to any of these questions, read this book to master a proven approach to building your enterprise data warehouse and BI initiatives. Extreme Scoping, based on the Business Intelligence Roadmap, will show you how to build analytics applications rapidly yet not sacrifice data management and enterprise architecture. In addition, all of the roles required to deliver all seven steps of this agile methodology are explained along with many real-world examples. From Wayne Eckerson's Foreword I've read many books about data warehousing and business intelligence (BI). This book by Larissa Moss is one of the best. I should not be surprised. Larissa has spent years refining the craft of designing, building, and delivering BI applications. Over the years, she has developed a keen insight about what works and doesn't work in BI. This book brings to light the wealth of that development experience. Best of all, this is not some dry text that laboriously steps readers through a technical methodology. Larissa expresses her ideas in a clear, concise, and persuasive manner. I highlighted so many beautifully written and insightful paragraphs in her manuscript that it became comical. I desperately wanted the final, published book rather than the manuscript so I could dog-ear it to death and place it front-and-center in my office bookshelf! From David Wellis Foreword Extreme Scoping is rich with advice and guidance for virtually every aspect of BI projects from planning and requirements to deployment and from back-end data management to front-end information and analytics services. Larissa is both a pragmatist and an independent thinker. Those qualities come through in the style of this book. Extreme Scoping is a well-written book that is easy to absorb. It is not full of surprises. It is filled with a lot of common sense and lessons learned through experience.

Data warehousing (DW) -- System development methodology (SDM) -- Agile system development methodology (ASDM) -- Extreme programming (XP) -- SCRUM -- Feature driven development (FDD) -- Agile data warehouse (ADW) methodology -- Data pakhuise -- Ontwikkelingstelselmethodologie -- Ratse ontwikkelingsstelselmethodologie -- Extremer programmering -- Funksiegedrewe ontwikkeling -- Ratse data pakhuise methodologie.

Introduction to Disciplined Agile Delivery 2nd Edition provides a quick overview of how agile software development works from beginning-to-end. It describes Disciplined Agile Delivery (DAD), the first of four levels of the Disciplined Agile (DA) process decision framework, and works through a case study describing a typical agile team's experiences adopting a DA approach. The book describes how the team develops the first release of a mission-critical application while working in a legacy enterprise environment. It describes their experiences from beginning-to-end, starting with their initial team initiation efforts through construction and finally to deploying the solution into production. It also describes how the team stays together for future releases, overseeing their process improvement efforts from their Scrum-based beginnings through to a lean continuous delivery approach that fits in with their organization's evolving DevOps strategy. The DAD framework is a hybrid of existing methods such as Scrum, Kanban, Agile Modeling, SAFe, Extreme Programming, Agile Data, Unified Process and many others. DAD provides the flexibility to use various approaches and plugs the gaps not addressed by mainstream agile methods. In a nutshell, DAD is "pragmatic agile." DAD describes proven strategies to adapt and scale your agile initiatives to suit the unique realities of your enterprise without having to figure it all out by yourself. Here's an overview of what each chapter covers: Chapter 1: Introduction. This chapter provides a quick overview of the book and a brief history of Disciplined Agile. Chapter 2: Reality over Rhetoric. This chapter explores several common myths about DAD and more importantly disproves them. Chapter 3: Disciplined Agile Delivery in a Nutshell. This chapter provides a brief yet comprehensive overview of DAD. Chapter 4: Introduction to the Case Study. This chapter introduces us to the team, describes the market opportunity that they hope to address, and describes the environment in which they're working. Chapter 5: Inception. The team's initiation effort includes initial requirements modeling and planning with their stakeholders in a streamlined manner, initial architecture modeling, setting up their physical work environment, setting up the start of their tooling infrastructure, initial risk identification, and finally securing stakeholder support and funding for the rest of the first release. Chapters 6 through 10: Construction. These chapters each describe a single Construction iteration, sharing the team's experiences during each of those two-week timeboxes. Chapter 11: Transition. The two-week transition phase focuses on final testing and fixing, training the support/help-desk staff, finishing a few short end-user "how to" videos, and deploying the solution into production. Chapter 12: The Road to Disciplined DevOps. This chapter overviews the team's improvement efforts over the next few releases, describing how they evolve from the agile Scrum-based lifecycle to a leaner approach and eventually to continuous delivery. All of this dovetails into their organization's efforts to implement a Disciplined DevOps strategy. Chapter 13: Closing Thoughts. This chapter overviews the disciplined agile resources that are available to you. Appendix: The Disciplined Agile Framework. This short appendix overviews our ongoing work on the Disciplined Agile framework to address the full scope of an agile business. At 111 pages, you should find this book to be a quick, informative read. What's Different in This Edition: Chapter 3 was completely rewritten to reflect the changes to DAD. Chapter 12 was rewritten to describe how the team evolved into a Disciplined DevOps strategy. Appendix A was rewritten to reflect the latest release of the DA framework. General updates were made throughout the book.

The Data Model Resource Book arms you with a set of proven data models and data warehouse designs for the core functions shared by most businesses. You get a comprehensive set of detailed models for marketing and sales, human resources, inventory, professional services, order processing, billing, product delivery, work order management, budgeting, accounting, and more. The authors also show you how to quickly convert the logical data models into enterprise-wide data warehouses as well as data marts.

Evolutionary Database Design (paperback)

Collaborative Dimensional Modeling, from Whiteboard to Star Schema

A Brief Guide to the Emerging World of Polyglot Persistence

The Agile Development of a Data Warehouse

A Comprehensive Guide for IT Professionals

A Disciplined Agile Delivery Handbook for Optimizing Your Way of Working

Agile Data Warehousing

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 book includes a selection of papers from the 2018 World Conference on Information Systems and Technologies (WorldCIST'18), held in Naples, Italy on March 27-29, 2018. WorldCIST is a global forum for researchers and practitioners to present and discuss recent results and innovations, current trends, professional experiences and the challenges of modern information systems and technologies research together with their technological development and applications. The main topics covered are: A) Information and Knowledge Management; B) Organizational Models and Information Systems; C) Software and Systems Modeling; D) Software Systems, Architectures, Applications and Tools; E) Multimedia Systems and Applications; F) Computer Networks, Mobility and Pervasive Systems; G) Intelligent and Decision Support Systems; H) Big Data Analytics and Applications; I) Human-Computer Interaction; J) Ethics, Computers & Security; K) Health Informatics; L) Information Technologies in Education; M) Information Technologies in Radiocommunications; N) Technologies for Biomedical Applications.

What can you do when database performance doesn't meet expectations? Before you turn to expensive hardware upgrades to solve the problem, reach for this book. Refactoring SQL Applications provides a set of tested options for making code modifications to dramatically improve the way your database applications function. Backed by real-world examples, you'll find quick fixes for simple problems, in-depth answers for more complex situations, and complete solutions for applications with extensive problems. Learn to: Determine if and where you can expect performance gains Apply quick fixes, such as limiting calls to the database in stored functions and procedures Refactor tasks, such as replacing application code by a stored procedure, or replacing iterative, procedural statements with sweeping SQL statements Refactor flow by increasing parallelism and switching business-induced processing from synchronous to asynchronous Refactor design using schema extensions, regular views, materialized views, partitioning, and more Compare before and after versions of a program to ensure you get the same results once you make modifications Refactoring SQL Applications teaches you to recognize and assess code that needs refactoring, and to understand the crucial link between refactoring and performance. If and when your application bogs down, this book will help you get it back up to speed.

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 Database Techniques

Architecture and Implementation

Delivering World-class Business Intelligence Systems Using Scrum and XP

Effective Strategies for the Agile Software Developer

The Definitive Guide to Dimensional Modeling

Corporate Information Factory