

Building A Scalable Data Warehouse With Data Vault 20

Thank you enormously much for downloading **Building A Scalable Data Warehouse With Data Vault 20** .Maybe you have knowledge that, people have look numerous time for their favorite books later than this Building A Scalable Data Warehouse With Data Vault 20 , but end taking place in harmful downloads.

Rather than enjoying a good book once a cup of coffee in the afternoon, on the other hand they juggled afterward some harmful virus inside their computer. **Building A Scalable Data Warehouse With Data Vault 20** is clear in our digital library an online right of entry to it is set as public in view of that you can download it instantly. Our digital library saves in multipart countries, allowing you to acquire the most less latency epoch to download any of our books in imitation of this one. Merely said, the Building A Scalable Data Warehouse With Data Vault 20 is universally compatible later than any devices to read.

An Introduction to Agile Data Engineering Using Data Vault 2.0 - Kent Graziano 2015-11-22

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.

Amazon Redshift Cookbook - Shruti Worlikar 2021-07-23

Discover how to build a cloud-based data warehouse at petabyte-scale that is burstable and built to scale for end-to-end analytical solutions Key FeaturesDiscover how to translate familiar data warehousing concepts into Redshift implementationUse impressive Redshift features to optimize development, productionizing, and operations processesFind out how to use advanced features such as concurrency scaling, Redshift Spectrum, and federated queriesBook Description Amazon Redshift is a fully managed, petabyte-scale AWS cloud data warehousing service. It enables you to build new data warehouse workloads on AWS and migrate on-premises traditional data warehousing platforms to Redshift. This book on Amazon Redshift starts by focusing on Redshift architecture, showing you how to perform database administration tasks on Redshift. You'll then learn how to optimize your data warehouse to quickly execute complex analytic queries against very large datasets. Because of the massive amount of data involved in data warehousing, designing your database for analytical processing lets you take full advantage of Redshift's columnar architecture and managed services. As you advance, you'll discover how to deploy fully automated and highly scalable extract, transform, and load (ETL) processes, which help minimize the operational efforts that you have to invest in managing regular ETL pipelines and ensure the timely and accurate refreshing of your data warehouse. Finally, you'll gain a clear understanding of Redshift use cases, data ingestion, data management, security, and scaling so that you can build a scalable data warehouse platform. By the end of this Redshift book, you'll be able to implement a Redshift-based data analytics solution and have understood the best practice solutions to commonly faced problems. What you will learnUse Amazon Redshift to build petabyte-scale data warehouses that are agile at scaleIntegrate your data warehousing solution with a data lake using purpose-built features and services on AWSBuild end-to-end analytical solutions from data sourcing to consumption with the help of useful recipesLeverage Redshift's comprehensive security capabilities to meet the most demanding business requirementsFocus on architectural insights and rationale when using analytical recipesDiscover best practices for working with big data to operate a fully managed solutionWho this book is for This book is for anyone involved in architecting, implementing, and optimizing an Amazon Redshift data warehouse, such as data warehouse developers, data analysts, database administrators, data engineers, and data scientists. Basic knowledge of data warehousing, database systems, and cloud concepts and familiarity with Redshift will be beneficial.

DW 2.0: The Architecture for the Next Generation of Data Warehousing - W.H. Inmon 2010-07-28

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

Super Charge Your Data Warehouse - Dan Linstedt 2011-11-01

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 Laney, 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"?

Pro SQL Server 2012 BI Solutions - Randal Root 2012-09-17

Business intelligence projects do not need to cost multi-millions of dollars or take months or even years to complete! Using rapid application development (RAD) techniques along with Microsoft SQL Server 2012, this book guides database administrators, SQL programmers, and report specialists in creating practical, cost-effective business intelligence solutions for their companies and departments. Pro SQL Server 2012 BI Solutions provides practical examples of cost-

effective business intelligence projects. Readers will be guided through several complete projects that build a foundation for real-world solutions. Even with limited experience using Microsoft's SQL Server, Integration Server, Analysis Server, and Reporting Server, you can leverage your existing knowledge of SQL programming and database design to provide users with the business intelligence reports they need. Provides recipes for multiple business intelligence scenarios Progresses from simple to advanced projects using several examples Shows Microsoft SQL Server technology used to complete real-world business intelligence projects
Data Architecture: A Primer for the Data Scientist - W.H. Inmon 2014-11-26

Today, the world is trying to create and educate data scientists because of the phenomenon of Big Data. And everyone is looking deeply into this technology. But no one is looking at the larger architectural picture of how Big Data needs to fit within the existing systems (data warehousing systems). Taking a look at the larger picture into which Big Data fits gives the data scientist the necessary context for how pieces of the puzzle should fit together. Most references on Big Data look at only one tiny part of a much larger whole. Until data gathered can be put into an existing framework or architecture it can't be used to its full potential. Data Architecture a Primer for the Data Scientist addresses the larger architectural picture of how Big Data fits with the existing information infrastructure, an essential topic for the data scientist. Drawing upon years of practical experience and using numerous examples and an easy to understand framework. W.H. Inmon, and Daniel Linstedt define the importance of data architecture and how it can be used effectively to harness big data within existing systems. You'll be able to: Turn textual information into a form that can be analyzed by standard tools. Make the connection between analytics and Big Data Understand how Big Data fits within an existing systems environment Conduct analytics on repetitive and non-repetitive data Discusses the value in Big Data that is often overlooked, non-repetitive data, and why there is significant business value in using it Shows how to turn textual information into a form that can be analyzed by standard tools Explains how Big Data fits within an existing systems environment Presents new opportunities that are afforded by the advent of Big Data Demystifies the murky waters of repetitive and non-repetitive data in Big Data

The Unified Star Schema: An Agile and Resilient Approach to Data Warehouse and Analytics Design - Bill Inmon 2020-10-03

Master the most agile and resilient design for building analytics applications: the Unified Star Schema (USS) approach. The USS has many benefits over traditional dimensional modeling. Witness the power of the USS as a single star schema that serves as a foundation for all present and future business requirements of your organization. Data warehouse legend Bill Inmon and business intelligence innovator, Francesco Puppini, explain step-by-step why the Unified Star Schema is the recommended approach for business intelligence designs today, and show through many examples how to build and use this new solution. This book contains two parts. Part I, Architecture, explains the benefits of data marts and data warehouses, covering how organizations progressed to their current state of analytics, and to the challenges that result from current business intelligence architectures. Chapter 1 covers the drivers behind and the characteristics of the data warehouse and data mart. Chapter 2 introduces dimensional modeling concepts, including fact tables, dimensions, star joins, and snowflakes. Chapter 3 recalls the evolution of the data mart. Chapter 4 explains Extract, Transform, and Load (ETL), and the value ETL brings to reporting. Chapter 5 explores the Integrated Data Mart Approach, and Chapter 6 explains how to monitor this environment. Chapter 7 describes the different types of metadata within the data warehouse environment. Chapter 8 progresses through the evolution to our current modern data warehouse environment. Part II, the Unified Star Schema, covers the Unified Star Schema (USS) approach and how it solves the challenges introduced in Part I. There are eight chapters within Part II: · Chapter 9, Introduction to the Unified Star Schema: Learn about its architecture and use cases, as well as how the USS approach differs from the traditional approach. · Chapter 10, Loss of Data: Learn about the loss of data and the USS Bridge. Understand that the USS approach does not create any join, and for this reason, it has no loss of data. · Chapter 11, The Fan Trap: Get introduced to the Oriented Data Model convention, and learn the dangers of a fan trap through an example. Differentiate join and association, and realize that an "in-memory association" is the preferred solution to the fan trap. · Chapter 12, The Chasm Trap: Become familiar with the Cartesian product, and then follow along with an example based on LinkedIn, which illustrates that a chasm trap

produces unwanted duplicates. See that the USS Bridge is based on a union, which does not create any duplicates. · Chapter 13, Multi-Fact Queries: Distinguish between multiple facts "with direct connection" versus multiple facts "with no direct connection". Explore how BI tools are capable of building aggregated virtual rows. · Chapter 14, Loops: Learn more about loops and five traditional techniques to solve them. Follow along with an implementation, which will illustrate the solution based on the USS approach. · Chapter 15, Non-Conformed Granularities: Learn about non-conformed granularities, and learn that the Unified Star Schema introduces a solution called "re-normalization". · Chapter 16, Northwind Case Study. Witness how easy it is to detect the pitfalls of Northwind using the ODM convention. Follow along with an implementation of the USS approach on the Northwind database with various BI tools.

Building a Scalable Data Warehouse with Data Vault 2.0 - Dan Linstedt 2015-09-15

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

Cassandra: The Definitive Guide - Jeff Carpenter 2016-06-29

Imagine what you could do if scalability wasn't a problem. With this hands-on guide, you'll learn how the Cassandra database management system handles hundreds of terabytes of data while remaining highly available across multiple data centers. This expanded second edition—updated for Cassandra 3.0—provides the technical details and practical examples you need to put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's non-relational design, with special attention to data modeling. If you're a developer, DBA, or application architect looking to solve a database scaling issue or future-proof your application, this guide helps you harness Cassandra's speed and flexibility. Understand Cassandra's distributed and decentralized structure Use the Cassandra Query Language (CQL) and cqlsh—the CQL shell Create a working data model and compare it with an equivalent relational model Develop sample applications using client drivers for languages including Java, Python, and Node.js Explore cluster topology and learn how nodes exchange data Maintain a high level of performance in your cluster Deploy Cassandra on site, in the Cloud, or with Docker Integrate Cassandra with Spark, Hadoop, Elasticsearch, Solr, and Lucene
The Data Warehouse Toolkit - Ralph Kimball 2011-08-08

This old edition was published in 2002. The current and final edition of this book is The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling, 3rd Edition which was published in 2013 under ISBN: 9781118530801. The authors begin with fundamental design recommendations and gradually progress step-by-step through increasingly complex scenarios. Clear-cut guidelines for designing dimensional models are illustrated using real-world data warehouse case studies drawn from a variety of business application areas and industries, including: Retail sales and e-commerce Inventory management Procurement Order management Customer relationship

management (CRM) Human resources management Accounting Financial services Telecommunications and utilities Education Transportation Health care and insurance By the end of the book, you will have mastered the full range of powerful techniques for designing dimensional databases that are easy to understand and provide fast query response. You will also learn how to create an architected framework that integrates the distributed data warehouse using standardized dimensions and facts.

Data Warehouse Systems - Alejandro Vaisman 2022-08-16

With this textbook, Vaisman and Zimányi deliver excellent coverage of data warehousing and business intelligence technologies ranging from the most basic principles to recent findings and applications. To this end, their work is structured into three parts. Part I describes “Fundamental Concepts” including conceptual and logical data warehouse design, as well as querying using MDX, DAX and SQL/OLAP. This part also covers data analytics using Power BI and Analysis Services. Part II details “Implementation and Deployment,” including physical design, ETL and data warehouse design methodologies. Part III covers “Advanced Topics” and it is almost completely new in this second edition. This part includes chapters with an in-depth coverage of temporal, spatial, and mobility data warehousing. Graph data warehouses are also covered in detail using Neo4j. The last chapter extensively studies big data management and the usage of Hadoop, Spark, distributed, in-memory, columnar, NoSQL and NewSQL database systems, and data lakes in the context of analytical data processing. As a key characteristic of the book, most of the topics are presented and illustrated using application tools.

Specifically, a case study based on the well-known Northwind database illustrates how the concepts presented in the book can be implemented using Microsoft Analysis Services and Power BI. All chapters have been revised and updated to the latest versions of the software tools used. KPIs and Dashboards are now also developed using DAX and Power BI, and the chapter on ETL has been expanded with the implementation of ETL processes in PostgreSQL. Review questions and exercises complement each chapter to support comprehensive student learning. Supplemental material to assist instructors using this book as a course text is available online and includes electronic versions of the figures, solutions to all exercises, and a set of slides accompanying each chapter. Overall, students, practitioners and researchers alike will find this book the most comprehensive reference work on data warehouses, with key topics described in a clear and educational style. “I can only invite you to dive into the contents of the book, feeling certain that once you have completed its reading (or maybe, targeted parts of it), you will join me in expressing our gratitude to Alejandro and Esteban, for providing such a comprehensive textbook for the field of data warehousing in the first place, and for keeping it up to date with the recent developments, in this current second edition.” From the foreword by Panos Vassiliadis, University of Ioannina, Greece.

Engineering the Transformation of the Enterprise - Stephan Aier 2021-09-01

The topics in this book cover a broad range of research interests: from business engineering and its application in corporate and business networking contexts to design science research as well as applied topics, where those research methods have been employed for modeling, data warehousing, information systems management, enterprise architecture management, management of large and complex projects, and enterprise transformation. The book is a Festschrift for Robert Winter in order to appreciate his work and to honor him as a personality with a high reputation in the information systems community. To this end, many professional colleagues or long-time companions both from the Institute of Information Management at the University of St. Gallen as well as from the international research community dedicated articles on topics related to Robert’s research. They reflect his ambition to uncompromisingly conduct high-class research that fuels the research community and at the same time contributes to improved industrial practice. The book is organized in three major parts: Part I “Business Engineering and Beyond” focuses on the methodology strongly shaped by Robert in St. Gallen with a focus on research being applied in corporate contexts. Part II “Design Science Research” spans from reflections on the practice of design science research to perspectives on design science research methodologies and eventually up to considerations to teach design science research methodology. Part III “Applied Fields” combines various applications of design science and related research methodologies with practical problems and future research topics.

Handbook of Research on Big Data Storage and Visualization Techniques - Segall, Richard S. 2018-01-05

The digital age has presented an exponential growth in the amount of data available to individuals looking to draw conclusions based on given or collected information across industries. Challenges associated with the analysis, security, sharing, storage, and visualization of large and complex data sets continue to plague data scientists and analysts alike as traditional data processing applications struggle to adequately manage big data. The Handbook of Research on Big Data Storage and Visualization Techniques is a critical scholarly resource that explores big data analytics and technologies and their role in developing a broad understanding of issues pertaining to the use of big data in multidisciplinary fields. Featuring coverage on a broad range of topics, such as architecture patterns, programing systems, and computational energy, this publication is geared towards professionals, researchers, and students seeking current research and application topics on the subject.

Rise of the Data Cloud - Frank Sloodman 2020-12-18

The rise of the Data Cloud is ushering in a new era of computing. The world’s digital data is mass migrating to the cloud, where it can be more effectively integrated, managed, and mobilized. The data cloud eliminates data siloes and enables data sharing with business partners, capitalizing on data network effects. It democratizes data analytics, making the most sophisticated data science tools accessible to organizations of all sizes. Data exchanges enable businesses to discover, explore, and easily purchase or sell data—opening up new revenue streams. Business leaders have long dreamed of data driving their organizations. Now, thanks to the Data Cloud, nothing stands in their way.

Google BigQuery: The Definitive Guide - Valliappa Lakshmanan 2019-10-23

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.

Data Warehousing in the Age of Big Data - Krish Krishnan 2013-05-02
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

Hands-On Data Warehousing with Azure Data Factory - Christian Coté 2018-05-31

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.

Building a Data Warehouse - Vincent Rainardi 2008-03-11

Here is the ideal field guide for data warehousing implementation. This book first teaches you how to build a data warehouse, including defining the architecture, understanding the methodology, gathering the requirements, designing the data models, and creating the databases. Coverage then explains how to populate the data warehouse and explores how to present data to users using reports and multidimensional databases and how to use the data in the data warehouse for business intelligence, customer relationship management, and other purposes. It also details testing and how to administer data warehouse operation.

Corporate Information Factory - W. H. Inmon 2002-03-14

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

Essential PySpark for Scalable Data Analytics - Sreeram Nudurupati 2021-10-29

Get started with distributed computing using PySpark, a single unified framework to solve end-to-end data analytics at scale Key Features Discover how to convert huge amounts of raw data into meaningful and actionable insights Use Spark's unified analytics engine for end-to-end analytics, from data preparation to predictive analytics Perform data ingestion, cleansing, and integration for ML, data analytics, and data visualization Book Description Apache Spark is a unified data analytics engine designed to process huge volumes of data quickly and efficiently. PySpark is Apache Spark's Python language API, which offers Python developers an easy-to-use scalable data analytics framework. Essential PySpark for Scalable Data Analytics starts by exploring the distributed computing paradigm and provides a high-level overview of Apache Spark. You'll begin your analytics journey with the data engineering process, learning how to perform data ingestion, cleansing, and integration at scale. This book helps you build real-time analytics pipelines that help you gain insights faster. You'll then discover methods for building cloud-based data lakes, and explore Delta Lake, which brings reliability to data lakes. The book also covers Data Lakehouse, an emerging paradigm, which combines the structure and performance of a data warehouse with the scalability of cloud-based data lakes. Later, you'll perform scalable data science and machine learning tasks using PySpark, such as data preparation, feature engineering, and model

training and productionization. Finally, you'll learn ways to scale out standard Python ML libraries along with a new pandas API on top of PySpark called Koalas. By the end of this PySpark book, you'll be able to harness the power of PySpark to solve business problems. What you will learn Understand the role of distributed computing in the world of big data Gain an appreciation for Apache Spark as the de facto go-to for big data processing Scale out your data analytics process using Apache Spark Build data pipelines using data lakes, and perform data visualization with PySpark and Spark SQL Leverage the cloud to build truly scalable and real-time data analytics applications Explore the applications of data science and scalable machine learning with PySpark Integrate your clean and curated data with BI and SQL analysis tools Who this book is for This book is for practicing data engineers, data scientists, data analysts, and data enthusiasts who are already using data analytics to explore distributed and scalable data analytics. Basic to intermediate knowledge of the disciplines of data engineering, data science, and SQL analytics is expected. General proficiency in using any programming language, especially Python, and working knowledge of performing data analytics using frameworks such as pandas and SQL will help you to get the most out of this book.

Building a Scalable Data Warehouse with Data Vault 2.0 - Dan Linstedt 2015-10-13

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

Graph Algorithms - Mark Needham 2019-05-16

Discover how graph algorithms can help you leverage the relationships within your data to develop more intelligent solutions and enhance your machine learning models. You'll learn how graph analytics are uniquely suited to unfold complex structures and reveal difficult-to-find patterns lurking in your data. Whether you are trying to build dynamic network models or forecast real-world behavior, this book illustrates how graph algorithms deliver value—from finding vulnerabilities and bottlenecks to detecting communities and improving machine learning predictions. This practical book walks you through hands-on examples of how to use graph algorithms in Apache Spark and Neo4j—two of the most common choices for graph analytics. Also included: sample code and tips for over 20 practical graph algorithms that cover optimal pathfinding, importance through centrality, and community detection. Learn how graph analytics vary from conventional statistical analysis Understand how classic graph algorithms work, and how they are applied Get guidance on which algorithms to use for different types of questions Explore algorithm examples with working code and sample datasets from Spark and Neo4j See how connected feature extraction can increase machine learning accuracy and precision Walk through creating an ML workflow for link prediction combining Neo4j and Spark

Data Architecture: A Primer for the Data Scientist - W.H. Inmon 2019-04-30

Over the past 5 years, the concept of big data has matured, data science has grown exponentially, and data architecture has become a standard

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

Business Intelligence Guidebook - Rick Sherman 2014-11-04

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.

Building and Managing the Meta Data Repository - David Marco 2000

"This is the first book to tackle the subject of meta data in data warehousing, and the results are spectacular . . . David Marco has written about the subject in a way that is approachable, practical, and immediately useful. *Building and Managing the Meta Data Repository: A Full Lifecycle Guide* is an excellent resource for any IT professional." - Steve Murchie Group Product Manager, Microsoft Corporation. Meta data repositories can provide your company with tremendous value if they are used properly and if you understand what they can, and can't, do. Written by David Marco, the industry's leading authority on meta data and well-known columnist for DM Review, this book offers all the guidance you'll need for developing, deploying, and managing a meta data repository to gain a competitive advantage. After illustrating the fundamental concepts, Marco shows you how to use meta data to increase your company's revenue and decrease expenses. You'll find a comprehensive look at the major trends affecting the meta data industry, as well as steps on how to build a repository that is flexible enough to adapt to future changes. This vendor-neutral guide also includes complete coverage of meta data sources, standards, and architecture, and it explores the full gamut of practical implementation issues. Taking you step-by-step through the process of implementing a meta data repository, Marco shows you how to:

- Evaluate meta data tools
- Build the meta data project plan
- Design a custom meta data architecture
- Staff a repository team
- Implement data quality through meta data
- Create a physical meta data model
- Evaluate meta data delivery requirements

The CD-ROM includes:

- A sample implementation project plan
- A function and feature checklist of meta data tool requirements
- Several physical meta data models to support specific business functions

Visit our Web site at

www.wiley.com/compbooks/ Visit the companion Web site at www.wiley.com/compbooks/marco

Cloud Analytics with Microsoft Azure - Has Altaiar 2019-11-26

Leverage the power of Azure to get efficient data insights from your big data in real time. Key Features: Explore the basics of cloud analytics using Azure. Discover different ways to process and visualize your data easily. Learn to use Azure Synapse Analytics (formerly known as Azure SQL Data Warehouse) to derive real-time customer insights. Book Description: With data being generated at an exponential speed, organizations all over the world are migrating their infrastructure to the cloud. Application management becomes much easier when you use a cloud platform to build, manage, and deploy your services and applications. *Cloud Analytics with Microsoft Azure* covers all that you need to extract useful insights from your data. You'll explore the power of data with big data analytics, the Internet of Things (IoT), machine learning, artificial intelligence, and DataOps. You'll also delve into data analytics by studying use cases that focus on creating actionable insights from near-real-time data. As you advance, you'll learn to build an end-to-end analytics pipeline on the cloud with machine learning and deep learning concepts. By the end of this book, you'll have developed a solid understanding of data analytics with Azure and its practical implementation. What you will learn: Explore the concepts of modern data warehouses and data pipelines. Discover different design considerations while applying a cloud analytics solution. Design an end-to-end analytics pipeline on the cloud. Differentiate between structured, semi-structured, and unstructured data. Choose a cloud-based service for your data analytics solutions. Use Azure services to ingest, store and analyze data of any scale. Who this book is for: If you're planning to adopt the cloud analytics model for your business, this book will help you understand the design and business considerations that you must keep in mind. Though not necessary, a basic understanding of data analytics concepts such as data streaming, data types, the machine learning life cycle, and Docker containers will help you get the most out of the book.

Design Patterns - Erich Gamma 1995

Software -- Software Engineering.

Modeling the Agile Data Warehouse with Data Vault - Hans Hultgren 2012-11-16

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.

Agile Data Warehouse Design - Lawrence Corr 2011-04-16

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 BEAM², an agile approach to dimensional modeling, for improving communication between data warehouse designers, BI stakeholders and the whole DW/BI development team. BEAM² 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 (BEAM²)
- 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 BEAM² dimensional shorthand notation
- Solving difficult DW/BI performance and usability problems with proven dimensional design patterns

Lawrence Corr is a data warehouse designer and educator. As Principal of DecisionOne Consulting, he helps clients to review and simplify their data warehouse designs, and advises vendors on visual data modeling techniques. He

regularly teaches agile dimensional modeling courses worldwide and has taught dimensional DW/BI skills to thousands of students. Jim Stagnitto is a data warehouse and master data management architect specializing in the healthcare, financial services, and information service industries. He is the founder of the data warehousing and data mining consulting firm Llumino.

Designing Data-Intensive Applications - Martin Kleppmann 2017-03-16
Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

The Modern Data Warehouse in Azure - Matt How 2020-06-15
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

Building the Data Warehouse - W. H. Inmon 2002-10-01
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.

Snowflake Cookbook - Hamid Mahmood Qureshi 2021-02-25
Develop modern solutions with Snowflake's unique architecture and integration capabilities; process bulk and real-time data into a data lake; and leverage time travel, cloning, and data-sharing features to optimize data operations Key Features Build and scale modern data solutions using the all-in-one Snowflake platform Perform advanced cloud analytics for implementing big data and data science solutions Make quicker and better-informed business decisions by uncovering key insights from your data Book Description Snowflake is a unique cloud-based data warehousing platform built from scratch to perform data management on the cloud. This book introduces you to Snowflake's unique architecture, which places it at the forefront of cloud data warehouses. You'll explore the compute model available with Snowflake, and find out how Snowflake allows extensive scaling through the virtual warehouses. You will then learn how to configure a virtual warehouse for optimizing cost and performance. Moving on, you'll get to grips with the data ecosystem and discover how Snowflake integrates with other technologies for staging and loading data. As you progress through the chapters, you will leverage Snowflake's capabilities to process a series of SQL statements using tasks to build data pipelines and find out how you can create modern data solutions and pipelines designed to provide high performance and scalability. You will also get to grips with creating role hierarchies, adding custom roles, and setting default roles for users before covering advanced topics such as data sharing, cloning, and performance optimization. By the end of this Snowflake book, you will be well-versed in Snowflake's architecture for building modern analytical solutions and understand best practices for solving commonly faced problems using practical recipes. What you will learn Get to grips with data warehousing techniques aligned with Snowflake's cloud architecture Broaden your skills as a data warehouse designer to cover the Snowflake ecosystem Transfer skills from on-premise data warehousing to the Snowflake cloud analytics platform Optimize performance and costs associated with a Snowflake solution Stage data on object stores and load it into Snowflake Secure data and share it efficiently for access Manage transactions and extend Snowflake using stored procedures Extend cloud data applications using Spark Connector Who this book is for This book is for data warehouse developers, data analysts, database administrators, and anyone involved in designing, implementing, and optimizing a Snowflake data warehouse. Knowledge of data warehousing and database and cloud concepts will be useful. Basic familiarity with Snowflake is beneficial, but not necessary.

Data Warehousing Fundamentals - Paulraj Ponniah 2004-04-07
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.

The Data Vault Guru - Patrick Cuba 2020-10-06
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 Enterprise Big Data Lake - Alex Gorelik 2019-02-21

The data lake is a daring new approach for harnessing the power of big data technology and providing convenient self-service capabilities. But is it right for your company? This book is based on discussions with practitioners and executives from more than a hundred organizations, ranging from data-driven companies such as Google, LinkedIn, and Facebook, to governments and traditional corporate enterprises. You'll learn what a data lake is, why enterprises need one, and how to build one successfully with the best practices in this book. Alex Gorelik, CTO and founder of Waterline Data, explains why old systems and processes can no longer support data needs in the enterprise. Then, in a collection of essays about data lake implementation, you'll examine data lake initiatives, analytic projects, experiences, and best practices from data experts working in various industries. Get a succinct introduction to data warehousing, big data, and data science Learn various paths enterprises take to build a data lake Explore how to build a self-service model and best practices for providing analysts access to the data Use different methods for architecting your data lake Discover ways to implement a data lake from experts in different industries

Agile Data Warehousing for the Enterprise - Ralph Hughes 2015-09-19

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 Elephant in the Fridge - John Giles 2019-04-15

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.

Data Management at Scale - Pietheinh Strengholt 2020-07-29

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 Pietheinh 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

Building the Data Lakehouse - Bill Inmon 2021-10

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