

Cloud Based Data Warehousing With Ibm Dashdb Mpp And Intel

Barry Devlin's Cloud Data Warehousing Volume 1 Book Release - Full Video What is a Data Warehouse? Database vs Data Warehouse vs Data Lake | What is the Difference? Transform traditional Data Warehouse with cloud and Big Data A book for SAP BW Professionals, explore of Cloud Data Warehousing to DataAI Analytics Modern Data Warehousing with BigQuery (Next '19 Rewind) Cloud-based Data Warehousing and Data Lake SolutionsSQL Trail Seasonal Series Winter 2022 - Pt 1 Cloud Data Warehousing: It's Time to Rethink Your Data Architecture Azure Event Subscription With Azure Event Grid to Store in Azure Storage Queue | Azure Event Driven Data Warehousing With BigQuery: Best Practices (Cloud Next '19) Best Practices: Cloud-Based Data Warehousing Enterprise-ready Data Warehousing in the Cloud Cloud based Data Warehouse Development - Case Study DAS Webinar: Cloud-Based Data Warehousing - What's New and What Stays the Same Cloud native data warehousing with Kubernetes // DoK Talks #106 Integrating Your Cloud-Based Data Lake and Data Warehouse For AI \u0026 BI Workloads Future of Data Meetup: Data Exploration and Dashboard Creation in a Cloud Based Data Warehouse Reimagine Data Warehousing: How The Home Depot is Using BigQuery to Scale (Cloud Next '19) What is Snowflake? 8 Minute Demo The Cloud Analytics E-Book - Powered by Smartbridge The Data Warehouse Toolkit
Data Warehousing, Business Intelligence and Analytics
Snowflake Cookbook
Managed Data Analysis in the Google Cloud
Business Intelligence Demystified
Cloud Data Warehousing Fd, Snowflake Special Edition (Custom)
Hands-on Amazon Redshift for Data Warehousing
DW 2.0: The Architecture for the Next Generation of Data Warehousing
Delivering the Promise of Big Data and Data Science
Cloud-Based Big Data Analytics in Vehicular Ad-Hoc Networks
Google BigQuery: The Definitive Guide
BigQuery for Data Warehousing
Foundations, Innovations, and Applications
SAP HANA 2.0
Multimedia Forensics and Security
Hands-On Data Warehousing with Azure Data Factory
Build modern data warehouses on Microsoft Azure
Computational Collective Intelligence
Official Google Cloud Certified Professional Data Engineer Study Guide
Data Management in the Cloud
Understand and Clear All Your Doubts and Misconceptions About BI (English Edition)

Cloud Based Data Warehousing With Ibm Dashdb Mpp And Intel

OMB No. 4082271634658 edited by

DAVENPORT JACK

The Data Warehouse Toolkit Springer

Enter the fast-paced world of SAP HANA 2.0 with this introductory guide. Begin with an exploration of the technological backbone of SAP HANA as a database and platform. Then, step into key SAP HANA user roles and discover core capabilities for administration, application development, advanced analytics, security, data integration, and more. No matter how SAP HANA 2.0 fits into your business, this book is your starting point. In this book, you'll learn about: a. Technology Discover what makes an in-memory database platform. Learn about SAP HANA's journey from version 1.0 to 2.0, take a tour of your technology options, and walk through deployment scenarios and implementation requirements. b. Tools Unpack your SAP HANA toolkit. See essential tools in action, from SAP HANA cockpit and SAP HANA studio, to the SAP HANA Predictive Analytics Library and SAP HANA smart data integration. c. Key Roles Understand how to use SAP HANA as a developer, administrator, data scientist, data center architect, and more. Explore key tasks like backend programming with SQLScript, security setup with roles and authorizations, data integration with the SAP HANA Data Management Suite, and more. Highlights include: 1) Architecture 2) Administration 3) Application development 4) Analytics 5) Security 6) Data integration 7) Data architecture 8) Data center
Data Warehousing, Business Intelligence and Analytics Prentice Hall

It's time to get your head in the cloud! In today's business environment, more and more people are requesting cloud-based solutions to help solve their business challenges. So how can you not only anticipate your clients' needs but also keep ahead of the curve to ensure their goals stay on track? With the help of this accessible book, you'll get a clear sense of cloud computing and understand how to communicate the benefits, drawbacks, and options to your clients so they can make the best choices for their unique needs. Plus, case studies give you the opportunity to relate real-life examples of how the latest technologies are giving organizations worldwide the opportunity to thrive as supply chain solutions in the cloud. Demonstrates how improvements in forecasting, collaboration, and inventory optimization can lead to cost savings Explores why cloud computing is becoming increasingly important Takes a close look at the types of cloud computing Makes sense of demand-driven forecasting using Amazon's cloud Whether you work in management, business, or IT, this is the dog-eared reference you'll want to keep close by as you continue making sense of the cloud.

Snowflake Cookbook Elsevier

As it is with building a house, most of the work necessary to build a data warehouse is neither visible nor obvious when looking at the completed product. While it may be easy to plan for a data warehouse that incorporates all the right concepts, taking the steps needed to create a warehouse that is as functional and user-friendly as it is theoreti
SAP PRESS

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 Lambda architecture by combining batch ETL/ELT jobs using Azure Data Factory with streaming technologies such as EventHub and Azure Databricks. You will discover how to manage cleansing and transformation prior to serving your blended datasets up to queries arriving from Azure SQL Data Warehouse. And you will know how to build robust analytic solutions on 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
Managed Data Analysis in the Google Cloud Technics Publications
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

Business Intelligence Demystified John Wiley & Sons
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.

Cloud Data Warehousing Fd, Snowflake Special Edition (Custom) Springer Science & Business Media

Data warehousing is one of the hottest business topics, and there's more to understanding data warehousing technologies than you might think. Find out the basics of data warehousing and how it facilitates data mining and business intelligence with *Data Warehousing For Dummies, 2nd Edition*. Data is probably your company's most important asset, so your data warehouse should serve your needs. The fully updated Second Edition of *Data Warehousing For Dummies* helps you understand, develop, implement, and use data warehouses, and offers a sneak peek into their future. You'll learn to: Analyze top-down and bottom-up data warehouse designs Understand the structure and technologies of data warehouses, operational data stores, and data marts Choose your project team and apply best development practices to your data warehousing projects Implement a data warehouse, step by step, and involve end-users in the process Review and upgrade existing data storage to make it serve your needs Comprehend OLAP, column-wise databases, hardware assisted databases, and middleware Use data mining intelligently and find what you need Make informed choices about consultants and data warehousing products *Data Warehousing For Dummies, 2nd Edition* also shows you how to involve users in the testing process and gain valuable feedback, what it takes to successfully manage a data warehouse project, and how to tell if your project is on track. You'll find it's the most useful source of data on the topic!

Hands-on Amazon Redshift for Data Warehousing Apress

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

DW 2.0: The Architecture for the Next Generation of Data Warehousing Packt Publishing Ltd

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, programming systems, and computational energy, this publication is geared towards professionals, researchers, and students seeking current research and application topics on the subject.

DELIVERING THE PROMISE OF BIG DATA AND DATA SCIENCE

CRC Press

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

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.

GOOGLE BIGQUERY: THE DEFINITIVE GUIDE

Morgan Kaufmann

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 Features Discover how to translate familiar data warehousing concepts into Redshift implementation Use impressive Redshift features to optimize development, productionizing, and operations processes Find out how to use advanced features such as concurrency scaling, Redshift Spectrum, and federated queries Book 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 learn Use Amazon Redshift to build petabyte-scale data warehouses that are agile at scale Integrate your data warehousing solution with a data lake using purpose-built features and services on AWS Build end-to-end analytical solutions from data sourcing to consumption with the help of useful recipes Leverage Redshift's comprehensive security capabilities to meet the most demanding business requirements Focus on architectural insights and rationale when using analytical recipes Discover best practices for working with big data to operate a fully managed solution Who 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.

BigQuery for Data Warehousing Apress

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.

Foundations, Innovations, and Applications SAP PRESS

Create a data warehouse, complete with reporting and dashboards using Google's BigQuery technology. This book takes you from the basic concepts of data warehousing through the design, build, load, and maintenance phases. You will build capabilities to capture data from the operational environment, and then mine and analyze that data for insight into making your business more successful. You will gain practical knowledge about

how to use BigQuery to solve data challenges in your organization. BigQuery is a managed cloud platform from Google that provides enterprise data warehousing and reporting capabilities. Part I of this book shows you how to design and provision a data warehouse in the BigQuery platform. Part II teaches you how to load and stream your operational data into the warehouse to make it ready for analysis and reporting. Parts III and IV cover querying and maintaining, helping you keep your information relevant with other Google Cloud Platform services and advanced BigQuery. Part V takes reporting to the next level by showing you how to create dashboards to provide at-a-glance visual representations of your business situation. Part VI provides an introduction to data science with BigQuery, covering machine learning and Jupyter notebooks. What You Will Learn Design a data warehouse for your project or organization Load data from a variety of external and internal sources Integrate other Google Cloud Platform services for more complex workflows Maintain and scale your data warehouse as your organization grows Analyze, report, and create dashboards on the information in the warehouse Become familiar with machine learning techniques using BigQuery ML Who This Book Is For Developers who want to provide business users with fast, reliable, and insightful analysis from operational data, and data analysts interested in a cloud-based solution that avoids the pain of provisioning their own servers.

SAP HANA 2.0 Springer

This book presents recent applications and approaches as well as challenges in digital forensic science. One of the evolving challenges that is covered in the book is the cloud forensic analysis which applies the digital forensic science over the cloud computing paradigm for conducting either live or static investigations within the cloud environment. The book also covers the theme of multimedia forensics and watermarking in the area of information security. That includes highlights on intelligence techniques designed for detecting significant changes in image and video sequences. Moreover, the theme proposes recent robust and computationally efficient digital watermarking techniques. The last part of the book provides several digital forensics related applications, including areas such as evidence acquisition enhancement, evidence evaluation, cryptography, and finally, live investigation through the importance of reconstructing the botnet attack scenario to show the malicious activities and files as evidences to be presented in a court.

MULTIMEDIA FORENSICS AND SECURITY

John Wiley & Sons

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

Hands-On Data Warehousing with Azure Data Factory

AuthorHouse

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. [Build modern data warehouses on Microsoft Azure](#) Morgan & Claypool Publishers

The deployment of Data Warehouses as a business application has grown tremendously over the past decade. Data warehouses are today considered to be one of the key components of an organizations overall IT strategy and architecture. This is especially true in the current Knowledge based global economy. Innovation and creativity is the current buzzword as business enterprises struggle to retain their stranglehold and find new markets for their products or services. Data warehouses are being developed and deployed for all businesses irrespective of its size and nature. Foreseeing a huge growth potential major hardware and software vendors, across the world, have quickly developed products and services specifically targeting the data warehousing market. The objective of this book is to provide the reader with an insight to the world of Data Warehousing, in a lucid manner devoid of mathematical complications.

COMPUTATIONAL COLLECTIVE INTELLIGENCE

"O'Reilly Media, Inc."

Vehicular traffic congestion and accidents remain universal issues in today's world. Due to the continued growth in the use of

vehicles, optimizing traffic management operations is an immense challenge. To reduce the number of traffic accidents, improve the performance of transportation systems, enhance road safety, and protect the environment, vehicular ad-hoc networks have been introduced. Current developments in wireless communication, computing paradigms, big data, and cloud computing enable the enhancement of these networks, equipped with wireless communication capabilities and high-performance processing tools. Cloud-Based Big Data Analytics in Vehicular Ad-Hoc Networks is a pivotal reference source that provides vital research on cloud and data analytic applications in intelligent transportation systems. While highlighting topics such as location routing, accident detection, and data warehousing, this publication addresses future challenges in vehicular ad-hoc networks and presents viable solutions. This book is ideally designed for researchers, computer scientists, engineers, automobile industry professionals, IT practitioners, academicians, and students seeking current research on cloud computing models in vehicular networks.

Official Google Cloud Certified Professional Data Engineer Study Guide For Dummies

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

Related with Cloud Based Data Warehousing With Ibm Dashdb Mpp And Intel:

© [Cloud Based Data Warehousing With Ibm Dashdb Mpp And Intel Examen Practico De Manejo Citas](#)

© [Cloud Based Data Warehousing With Ibm Dashdb Mpp And Intel Example Of Unit Form In Math](#)

© [Cloud Based Data Warehousing With Ibm Dashdb Mpp And Intel Example Of Sociological Imagination](#)