OMB No. 6024597253811

Data Lake Development With Big Data

Database vs Data Warehouse vs Data Lake | What is the Difference? Journey to the Data Lake How Progressive Paved a Faster, Big Data In 5 Minutes | What Is Big Data?| Big Data Analytics | Big Data Tutorial | Simplilearn What is a Data Lake? 10 Data Books You Should Read in 2022 Azure Data Lake Storage (Gen 2) Tutorial | Best storage solution for big data analytics in Azure Back to Basics: Building an Efficient Data Lake History and Evolution of Data Lake Architecture - Post Lambda Architecture Data Lake Architecture Designing Data Lakes: Best Practices (Level 200) What is Azure Data Lake and When to Use It Serverless Architecture for Big Data, and Data Lake | XenonStack The Big Problem With Data Lake Architecture Data Lakes for Big Data MOOC from EMC Intro to Data Lakehouse Big Data that is easy and productive with Azure Data Lake Processing Big Data with Azure Data Lake Analytics | Microsoft on edX | Course About Video Data Lakehouses Explained Building a Data Lake on AWS Big Data Architectures and The Data Lake | James Serra | Cloud

Next Generation Databases
Data Engineering with Apache Spark, Delta Lake, and Lakehouse
Architecting in the Cloud with Azure Data Lake, HDInsight, and Spark
8th International Conference, BDA 2020, Sonepat, India, December 15–18, 2020, Proceedings
Modern Data Strategy
Delivering Data-Driven Value at Scale
Data Lakes For Dummies
The AI Ladder
SAP HANA 2.0
NoSQL Distilled
Managing Successful Data Projects
Handbook of Research on Pattern Engineering System Development for Big Data Analytics
Big Data Applications in Industry 4. 0
Conceptual Modeling Perspectives
Big Data in Practice
The Semantic Web: ESWC 2019 Satellite Events
ESWC 2019 Satellite Events, Portorož, Slovenia, June 2–6, 2019, Revised Selected Papers
A Brief Guide to the Emerging World of Polyglot Persistence
The Enterprise Big Data Lake
Building the Data Lakehouse
Big Data Integration

Data Lake Development With Big Data *OMB No. 6024597253811 edited by*

KEIRA ZACHARY

Next Generation Databases Apress
Explore architectural approaches to building Data Lakes that ingest, index, manage, and analyze massive amounts of data using Big Data technologies About This Book Comprehend the intricacies of architecting a Data Lake and build a data strategy around your current data architecture Efficiently manage vast amounts of data and deliver it to multiple applications and systems with a high degree of performance and scalability Packed with industry best practices and use-case scenarios to get you up-and-running Who This Book Is For This book is for architects and senior managers who are responsible for building a strategy around their current data architecture, helping them identify the need for a Data Lake implementation in an enterprise context. The reader will need a good knowledge of master data management and information lifecycle management, and experience of Big Data technologies. What You Will Learn Identify the need for a Data Lake in your enterprise context and learn to architect a Data Lake Learn to build various tiers of a Data Lake, such as data intake, management, consumption, and governance, with a focus on practical implementation scenarios Find out the key considerations to be taken into account while building each tier of the Data Lake Understand Hadoop-oriented data transfer mechanism to ingest data in batch, micro-batch, and real-time modes Explore various data integration needs and learn how to perform data enrichment and data transformations using Big Data technologies Enable data discovery on the Data Lake to allow users to discover the data Discover how data is packaged and provisioned for consumption Comprehend the importance of including data governance disciplines while building a Data Lake In Detail A Data Lake is a highly scalable platform for storing huge volumes of multistructured data from disparate sources with centralized data management services. This book explores the potential of Data Lakes and explores architectural approaches to building data lakes that ingest, index, manage, and analyze massive amounts of data using batch and real-time processing frameworks. It guides you on how to go about building a Data Lake that is managed by Hadoop and accessed as required by other Big Data applications. This book will guide readers (using best practices) in developing Data Lake's capabilities. It will focus on architect data governance, security, data quality, data lineage tracking, metadata management, and semantic data tagging. By the end of this book, you will have a good understanding of building a Data Lake for Big Data. Style and approach Data Lake Development with Big Data provides architectural approaches to building a Data Lake. It follows a use case-based approach where practical implementation scenarios of each key component are explained. It also helps you understand how these use cases are implemented in a Data Lake. The chapters are organized in a way that mimics the sequential data flow evidenced in a Data Lake.

Data Engineering with Apache Spark, Delta Lake, and Lakehouse AuthorHouse
While many companies ponder implementation details such as distributed processing engines and algorithms for data analysis, this practical book takes a much wider view of big data development, starting with initial planning and moving diligently toward execution. Authors Ted Malaska and Jonathan Seidman guide you through the major components necessary to start, architect, and develop successful big data projects. Everyone from CIOs and COOs to lead architects and developers will explore a variety of big data architectures and applications, from massive data pipelines to web-scale applications. Each chapter addresses a piece of the software development life cycle and identifies patterns to maximize long-term success throughout the life of your project. Start the planning process by considering the key data project types Use guidelines to evaluate and select data management solutions Reduce risk related to technology, your team, and vague requirements Explore system interface design using APIs, REST, and pub/sub systems Choose the right distributed storage system for your big data system Plan and implement metadata collections for your data architecture Use data pipelines to ensure data integrity from source to final storage Evaluate the attributes of various engines for processing the data you collect

Architecting in the Cloud with Azure Data Lake, HDInsight, and Spark

O'Reilly Media
Many enterprises are investing in a next-generation data lake, hoping to democratize data at scale to provide business insights and ultimately make automated intelligent decisions. In this practical book, author Zhamak Dehghani reveals that, despite the time, money, and effort poured into them, data warehouses and data lakes fail when applied at the scale and speed of today's organizations. A distributed data mesh is a better choice. Dehghani guides architects, technical leaders, and decision makers on their journey from monolithic big data architecture to a paradigm that draws from modern distributed architecture. A data mesh considers domains as a first-class concern, applies platform thinking to create self-serve data infrastructure, and treats data as a product. This book shows you why and how. Examine the current landscape of data architectures, their underlying characteristics, and failure modes Learn how to divide data (and its supporting technology stacks and architecture) into operational data and analytical data Get a complete introduction to data mesh principles and logical architecture Create a foundation for gaining value from analytical data and historical facts at scale Move beyond a monolithic data lake to a distributed data mesh

8th International Conference, BDA 2020, Sonepat, India, December 15–18, 2020, Proceedings IGI Global
Data pipelines are the foundation for success in data analytics. Moving data from numerous diverse sources and transforming it to provide context is the difference between having data and actually gaining value from it. This pocket reference defines data pipelines and explains how they work in today's modern data stack. You'll learn common considerations and key decision points when implementing pipelines, such as batch versus streaming data ingestion and build versus buy. This book addresses the most common decisions made by data professionals and discusses foundational concepts that apply to open source frameworks, commercial products, and homegrown solutions. You'll learn: What a data pipeline is and how it works How data is moved and processed on modern data infrastructure, including cloud platforms Common tools and products used by data engineers to build pipelines How pipelines support analytics and reporting needs Considerations for pipeline maintenance, testing, and alerting

Modern Data Strategy Springer
The best-selling author of Big Data is back, this time with a unique and in-depth insight into how specific companies use big data. Big data is on the tip of everyone's tongue. Everyone understands its power and importance, but many fail to grasp the actionable steps and resources required to utilise it effectively. This book fills the knowledge gap by showing how major companies are using big data every day, from an up-close, on-the-ground perspective. From technology, media and retail, to sport teams, government agencies and financial institutions, learn the actual strategies and processes being used to learn about customers, improve manufacturing, spur innovation, improve safety and so much more. Organised for easy dip-in navigation, each chapter follows the same structure to give you the information you need quickly. For each company profiled, learn what data was used, what problem it solved and the processes put it place to make it practical, as well as the technical details, challenges and lessons learned from each unique scenario. Learn how predictive analytics helps Amazon, Target, John Deere and Apple understand their customers Discover how big data is behind the success of Walmart, LinkedIn, Microsoft and more Learn how big data is changing medicine, law enforcement, hospitality, fashion, science and banking Develop your own big data strategy by accessing additional reading materials at the end of each chapter

Delivering Data-Driven Value at Scale McGraw Hill Professional
Conceptual modeling has always been one of the main issues in information systems engineering as it aims to describe the general knowledge of the system at an abstract level that facilitates user understanding and software development. This collection of selected papers provides a comprehensive and extremely readable overview of what conceptual modeling is and perspectives on making it more and more relevant in our society. It covers topics like modeling the human genome, blockchain technology, model-driven software development, data integration, and wiki-like repositories and demonstrates the general applicability of conceptual modeling to various problems in diverse domains. Overall, this book is a source of inspiration for everybody in academia working on the vision of creating a strong, fruitful and creative community of conceptual modelers. With this book the editors and authors want to honor Prof. Antoni Olivé for his enormous and ongoing contributions to the conceptual modeling discipline. It was presented to him on the occasion of his keynote at ER 2017 in Valencia, a conference that he has contributed to and supported for over 20 years. Thank you very much to Antoni for so many years of cooperation and friendship.

Data Lakes For Dummies John Wiley & Sons
Organizations invest incredible amounts of time and money obtaining and then storing big data in data stores called data lakes. But how many of these organizations can actually get the data back out in a useable form? Very few can turn the data lake into an information gold mine. Most wind up with garbage dumps. Data Lake Architecture will explain how to build a useful data lake, where data scientists and data analysts can solve business challenges and identify new business opportunities. Learn how to structure data lakes as well as analog, application, and text-based data ponds to

provide maximum business value. Understand the role of the raw data pond and when to use an archival data pond. Leverage the four key ingredients for data lake success: metadata, integration mapping, context, and metaprocess. Bill Inmon opened our eyes to the architecture and benefits of a data warehouse, and now he takes us to the next level of data lake architecture.
The AI Ladder "O'Reilly Media, Inc."
Get a 360-degree view of how the journey of data analytics solutions has evolved from monolithic data stores and enterprise data warehouses to data lakes and modern data warehouses. You will This book includes comprehensive coverage of how: To architect data lake analytics solutions by choosing suitable technologies available on Microsoft Azure The advent of microservices applications covering ecommerce or modern solutions built on IoT and how real-time streaming data has completely disrupted this ecosystem These data analytics solutions have been transformed from solely understanding the trends from historical data to building predictions by infusing machine learning technologies into the solutions Data platform professionals who have been working on relational data stores, non-relational data stores, and big data technologies will find the content in this book useful. The book also can help you start your journey into the data engineer world as it provides an overview of advanced data analytics and touches on data science concepts and various artificial intelligence and machine learning technologies available on Microsoft Azure. What Will You Learn You will understand the: Concepts of data lake analytics, the modern data warehouse, and advanced data analytics Architecture patterns of the modern data warehouse and advanced data analytics solutions Phases—such as Data Ingestion, Store, Prep and Train, and Model and Serve—of data analytics solutions and technology choices available on Azure under each phase In-depth coverage of real-time and batch mode data analytics solutions architecture Various managed services available on Azure such as Synapse analytics, event hubs, Stream analytics, CosmosDB, and managed Hadoop services such as Databricks and HDInsight Who This Book Is For Data platform professionals, database architects, engineers, and solution architects

SAP HANA 2.0

Technics Publications
A Dell Technologies perspective on today's data landscape and the key ingredients for planning a modern, distributed data pipeline for your multicloud data-driven enterprise
NoSQL Distilled IBM Redbooks
"It’s not easy to find such a generous book on big data and databases. Fortunately, this book is the one." Feng Yu. Computing Reviews. June 28, 2016. This is a book for enterprise architects, database administrators, and developers who need to understand the latest developments in database technologies. It is the book to help you choose the correct database technology at a time when concepts such as Big Data, NoSQL and NewSQL are making what used to be an easy choice into a complex decision with significant implications. The relational database (RDBMS) model completely dominated database technology for over 20 years. Today this "one size fits all" stability has been disrupted by a relatively recent explosion of new database technologies. These paradigm-busting technologies are powering the "Big Data" and "NoSQL" revolutions, as well as forcing fundamental changes in databases across the board. Deciding to use a relational database was once truly a no-brainer, and the various commercial relational databases competed on price, performance, reliability, and ease of use rather than on fundamental architectures. Today we are faced with choices between radically different database technologies. Choosing the right database today is a complex undertaking, with serious economic and technological consequences. Next Generation Databases demystifies today’s new database technologies. The book describes what each technology was designed to solve. It shows how each technology can be used to solve real word application and business problems. Most importantly, this book highlights the architectural differences between technologies that are the critical factors to consider when choosing a database platform for new and upcoming projects. Introduces the new technologies that have revolutionized the database landscape Describes how each technology can be used to solve specific application or business challenges Reviews the most popular new wave databases and how they use these new database technologies
Managing Successful Data Projects CRC Press
This IBM RedguideTM publication looks back on the key decisions that made the data lake successful and looks forward to the future. It proposes that the metadata management and governance approaches developed for the data lake can be adopted more broadly to increase the value that an organization gets from its data. Delivering this broader vision, however, requires a new generation of data catalogs and governance tools built on open standards that are adopted by a multi-vendor ecosystem of data platforms and tools. Work is already underway to define and deliver this capability, and there are multiple ways to engage. This guide covers the reasons why this new capability is critical for modern businesses and how you can get value from it.
Handbook of Research on Pattern Engineering System Development for Big Data Analytics Apress
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 learnhow 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 Piethein 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
Big Data Applications in Industry 4. 0 IBM Redbooks
The need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational “NoSQL” databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage's pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Polyglot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access.

Conceptual Modeling Perspectives

"O'Reilly Media, Inc."
The big data era is upon us: data are being generated, analyzed, and used at an unprecedented scale, and data-driven decision making is sweeping through all aspects of society. Since the value of data explodes when it can be linked and fused with other data, addressing the big data integration (BDI) challenge is critical to realizing the promise of big data. BDI differs from traditional data integration along the dimensions of volume, velocity, variety, and veracity. First, not only can data sources contain a huge volume of data, but also the number of data sources is now in the millions. Second, because of the rate at which newly collected data are made available, many of the data sources are very dynamic, and the number of data sources is also rapidly exploding. Third, data sources are extremely heterogeneous in their structure and content, exhibiting considerable variety even for substantially similar entities. Fourth, the data sources are of widely differing qualities, with significant differences in the coverage, accuracy and timeliness of data provided. This book explores the progress that has been made by the data integration community on the topics of schema alignment, record linkage and data fusion in addressing these novel challenges faced by big data integration. Each of these topics is covered in a systematic way: first starting with a quick tour of the topic in the context of traditional data integration, followed by a detailed, example-driven exposition of recent innovative techniques that have been proposed to address the BDI challenges of volume, velocity, variety, and veracity. Finally, it presents merging topics and opportunities that are specific to BDI, identifying promising directions for the data integration community.
Big Data in Practice John Wiley & Sons
The Microsoft Azure cloud is an ideal platform for data-intensive applications. Designed for productivity, Azure provides pre-built services that make collection, storage, and analysis much easier to implement and manage. Azure Storage, Streaming, and Batch Analytics teaches you how to design a reliable, performant, and cost-effective data infrastructure in Azure by progressively building a complete working analytics system. Summary The Microsoft Azure cloud is an ideal platform for data-intensive applications. Designed for productivity, Azure provides pre-built services that make collection, storage, and analysis much easier to implement and manage. Azure Storage, Streaming, and Batch Analytics teaches you how to design a reliable, performant, and cost-effective data infrastructure in Azure by progressively building a complete working analytics system. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Microsoft Azure provides dozens of services that simplify storing and processing data. These services are secure, reliable, scalable, and cost efficient. About the book Azure Storage, Streaming, and Batch Analytics shows you how to build state-of-the-art data solutions with tools from the Microsoft Azure platform. Read along to construct a cloud-native data warehouse, adding features like real-time data processing. Based on the Lambda architecture for big data, the design uses scalable services such as Event Hubs, Stream Analytics, and SQL databases. Along the way, you’ll cover most of the topics needed to earn an Azure data engineering certification. What's inside Configuring Azure services for speed and cost Constructing data pipelines with Data Factory Choosing the right data storage methods About the reader For readers familiar with database management. Examples in C# and PowerShell. About the author Richard Nuckolls is a senior developer building big data analytics and reporting systems in Azure. Table of Contents 1 What is data engineering? 2 Building an analytics system in Azure 3 General storage with Azure Storage accounts 4 Azure Data Lake Storage 5 Message handling with Event Hubs 6 Real-time queries with Azure Stream Analytics 7 Batch queries with Azure Data Lake Analytics 8 U-SQL for complex analytics 9 Integrating with Azure Data Lake Analytics 10 Service integration with Azure Data Factory 11 Managed SQL with Azure SQL Database 12 Integrating Data Factory with SQL Database 13 Where to go next
The Semantic Web: ESWC 2019 Satellite Events "O'Reilly Media, Inc."
Helps users understand the breadth of Azure services by organizing them into a reference framework they can use when crafting their own big-data analytics solution.

ESWC 2019 Satellite Events, Portoro□, Slovenia, June 2–6, 2019, Revised Selected Papers

"O'Reilly Media, Inc."
This book constitutes the proceedings of the 8th International Conference on Big Data Analytics, BDA 2020, which took place during December 15-18, 2020, in Sonepat, India. The 11 full and 3 short papers included in this volume were carefully reviewed and selected from 48 submissions; the book also contains 4 invited and 3 tutorial papers. The contributions were organized in topical sections named as follows: data science systems; data science architectures; big data analytics in healthcare; information interchange of Web data resources; and business analytics.

A Brief Guide to the Emerging World of Polyglot Persistence

Morgan Kaufmann
This book constitutes the thoroughly refereed post-conference proceedings of the Satellite Events of the 16th Extended Semantic Web Conference, ESWC 2019, held in Portorož, Slovenia, in June 2019. The volume contains 38 poster and demonstration papers, 2 workshop papers,5 PhD symposium papers, and 3 industry track papers, selected out of a total of 68 submissions. They deal with all areas of semantic web research, semantic technologies on the Web and Linked Data.
The Enterprise Big Data Lake Apress
Perform fast interactive analytics against different data sources using the Trino high-performance distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Trino. Initially developed by Facebook, open source Trino is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber, and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Trino query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Trino's use cases and learn about tools that will help you connect to Trino and query data Go deeper: Learn Trino's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Trino in production: Secure Trino, monitor workloads, tune queries, and connect more applications; learn how other organizations apply Trino
Building the Data Lakehouse John Wiley & Sons
Software Architecture for Big Data and the Cloud is designed to be a single resource that brings together research on how software architectures can solve the challenges imposed by building big data software systems. The challenges of big data on the software architecture can relate to scale, security, integrity, performance, concurrency, parallelism, and dependability, amongst others. Big data handling requires rethinking architectural solutions to meet functional and non-functional requirements related to volume, variety and velocity. The book's editors have varied and complementary backgrounds in requirements and architecture, specifically in software architectures

for cloud and big data, as well as expertise in software engineering for cloud and big data. This book brings together work across different disciplines in software engineering, including work expanded from conference tracks and workshops led by the editors. Discusses systematic and disciplined approaches to building software architectures for cloud and big data with state-of-the-art methods and techniques Presents case studies involving enterprise, business, and government service deployment of big data applications Shares guidance on theory, frameworks, methodologies, and architecture for cloud and big data