1. Record Nr. UNINA9910799251003321

Autore	Kumar Anjani
Titolo	Architecting a Modern Data Warehouse for Large Enterprises : Build Multi-Cloud Modern Distributed Data Warehouses with Azure and AWS / / Anjani Kumar, Abhishek Mishra, and Sanjeev Kumar
Pubbl/distr/stampa	New York, NY : , : Apress Media LLC, , [2024] ©2024
ISBN	979-88-6880-029-0
Edizione	[First edition.]
Descrizione fisica	1 online resource (XV, 368 p. 146 illus.)
Disciplina	005.7
Soggetti	Data warehousing Microsoft Azure (Computing platform) Big data Business enterprises - Data processing
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Chapter 1: Introduction -- Chapter 2: Modern Data Warehouses -- Chapter 3: Data Lake, Lake House, and Delta Lake -- Chapter 4: Data Mesh -- Chapter 5: Data Orchestration Techniques -- Chapter 6: Data Democratization, Governance, and Security -- Chapter 7: Business Intelligence.
Sommario/riassunto	Design and architect new generation cloud-based data warehouses using Azure and AWS. This book provides an in-depth understanding of how to build modern cloud-native data warehouses, as well as their history and evolution. The book starts by covering foundational data warehouse concepts, and introduces modern features such as distributed processing, big data storage, data streaming, and processing data on the cloud. You will gain an understanding of the synergy, relevance, and usage data warehousing standard practices in the modern world of distributed data processing. The authors walk you through the essential concepts of Data Mesh, Data Lake, Lakehouse, and Delta Lake. And they demonstrate the services and offerings available on Azure and AWS that deal with data orchestration, data democratization, data governance, data security, and business

intelligence. After completing this book, you will be ready to design and architect enterprise-grade, cloud-based modern data warehouses using industry best practices and guidelines. You will: Understand the core concepts underlying modern data warehouses Design and build cloud-native data warehouses Gain a practical approach to architecting and building data warehouses on Azure and AWS Implement modern data warehousing components such as Data Mesh, Data Lake, Delta Lake, and Lakehouse Process data through pandas and evaluate your model’s performance using metrics such as F1-score, precision, and recall Apply deep learning to supervised, semi-supervised, and unsupervised anomaly detection tasks for tabular datasets and time series applications.