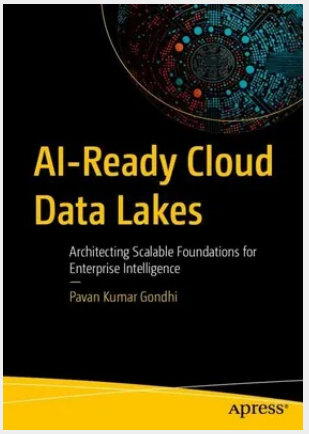
Gondhi

AI-Ready Cloud Data Lakes

Architecting Scalable Foundations for Enterprise Intelligence



This book provides a comprehensive guide to designing, building, governing, and operating modern cloud data lakes in the era of AI and enterprise-scale analytics. Moving beyond vendor marketing and surface-level cloud tutorials, it focuses on the challenges that actually emerge at scale. Readers will learn how to evaluate and implement modern data architectures, compare data lakes, warehouses, and lakehouses, and make informed decisions around open table formats including Apache Iceberg, Delta Lake, and Apache Hudi. The book also explores multi-cloud reference architectures across AWS, Azure, and Google Cloud Platform, helping organizations build scalable, resilient, and future-ready data platforms while minimizing vendor lock-in.
Covering the complete lifecycle from architecture and implementation to governance, security, operations, and AI integration, the book provides practical guidance on data ingestion, metadata management, Infrastructure-as-Code, observability, and cost optimization. It treats AI-readiness as a core architectural constraint rather than an afterthought, demonstrating how to build platforms that support machine learning, vector embeddings, retrieval-augmented generation (RAG), and generative AI workloads from day one. The governance and compliance chapters are grounded in real-world enterprise requirements, mapping best practices to frameworks such as BCBS 239, SR 11-7, FFIEC AI guidance, and EU AI Act Article 10. Through migration playbooks, implementation blueprints, and lessons learned from production environments, readers will gain a practical roadmap for building secure, compliant, AI-ready cloud data lakes that deliver measurable business value and support enterprise intelligence at scale..
You Will:- Design scalable cloud-native data lake architectures using AWS, Azure, and Google Cloud platforms.- Evaluate data lakes, warehouses, and lakehouses to select the right architecture for your business and AI initiatives.- Implement modern data ingestion, metadata management, and open table formats such as Apache Iceberg, Delta Lake, and Apache Hudi.- Build AI-ready data platforms that support machine learning, vector embeddings, retrieval-augmented generation (RAG), and generative AI workloads.
This book is for : Senior data engineers, Enterprise architects, principal and staff engineers

This book provides a comprehensive guide to designing, building, governing, and operating modern cloud data lakes in the era of AI and enterprise-scale analytics. Moving beyond vendor marketing and surface-level cloud tutorials, it focuses on the challenges that actually emerge at scale. Readers will learn how to evaluate and implement modern data architectures, compare data lakes, warehouses, and lakehouses, and make informed decisions around open table formats including Apache Iceberg, Delta Lake, and Apache Hudi. The book also explores multi-cloud reference architectures across AWS, Azure, and Google Cloud Platform, helping organizations build scalable, resilient, and future-ready data platforms while minimizing vendor lock-in.
Covering the complete lifecycle from architecture and implementation to governance, security, operations, and AI integration, the book provides practical guidance on data ingestion, metadata management, Infrastructure-as-Code, observability, and cost optimization. It treats AI-readiness as a core architectural constraint rather than an afterthought, demonstrating how to build platforms that support machine learning, vector embeddings, retrieval-augmented generation (RAG), and generative AI workloads from day one. The governance and compliance chapters are grounded in real-world enterprise requirements, mapping best practices to frameworks such as BCBS 239, SR 11-7, FFIEC AI guidance, and EU AI Act Article 10. Through migration playbooks, implementation blueprints, and lessons learned from production environments, readers will gain a practical roadmap for building secure, compliant, AI-ready cloud data lakes that deliver measurable business value and support enterprise intelligence at scale..
You Will:- Design scalable cloud-native data lake architectures using AWS, Azure, and Google Cloud platforms.- Evaluate data lakes, warehouses, and lakehouses to select the

64,19 €
64,19 € (zzgl. MwSt.)

vorbestellbar, Erscheinungstermin ca. April 2027

Artikelnummer: 9798868834547
Medium: Buch
ISBN: 979-8-8688-3454-7
Verlag: APRESS L.P.
Erscheinungstermin: 14.04.2027
Sprache(n): Englisch
Auflage: 1. Auflage 2027
Produktform: Kartoniert
Format (B x H): 178 x 254 mm

Kundenservice Verlag Dr. Otto Schmidt KG
Neumannstraße 10, 40235 Düsseldorf | kundenservice@otto-schmidt.de | 0221-93738-997 (Inland)