

The Database Safari

Exploring Vector, Time-Series, Graph, Columnar/OLAP, and Geospatial Databases in the Wild

Pack your field kit—this is a guided safari through today's database ecosystem. In the wild, one size rarely fits all: Modern systems often need specialized databases to deliver real-world performance and scale. This book is a practical field guide to those less commonly used databases—vector, time-series, graph, columnar/OLAP, geospatial, event-stream, and more. It shows you how each engine works, where it excels, and how to choose the right engine for each workload. The book provides clear trade-off tables, indexing strategies, query patterns, and use-case recipes you can apply immediately.

Starting with foundations in SQL and NoSQL, this book explores each “species” chapter by chapter, emphasizing hands-on scenarios and pitfalls to avoid. It then brings everything together in polyglot/hybrid architectures, helping you combine engines confidently, whether you're designing a similarity search with vector databases (ANN), building observability pipelines on time-series data and event logs, modeling relationships in graph databases, or delivering analytics with columnar/OLAP and geospatial queries. The result is a concise, engineer-first map of the database landscape that turns theory into pragmatic decision-making.


What You Will Learn

- Identify the right engine for the job across vector, time-series, graph, columnar/OLAP, geospatial, and event-stream databases
- Evaluate trade-offs in performance, scalability, consistency, storage layout, indexing, and operational complexity
- Design polyglot architectures that combine multiple engines (e.g., vector + OLAP + event-stream) for end-to-end solutions
- Apply use-case patterns for similarity search (ANN), observability metrics, fraud/network analysis, geospatial routing, and analytical reporting
- Understand core internals that matter (index types, retention policies, log-structured storage, spatial indexes, columnar execution) without vendor lock-in
- Avoid common pitfalls in schema design, query planning, and cross-engine data movement
- Build a decision checklist and trade-off tables to justify engine selection to stakeholders

Who This Book Is For

Software engineers; back-end developers; data engineers; system architects; site reliability engineers (SREs); technical leads with basic SQL/NoSQL experience who need practical guidance to select, compare, and combine specialized databases for real-world applications

Pack your field kit—this is a guided safari through today's database ecosystem. In the wild, one size rarely fits all: Modern systems often need specialized databases to deliver real-world performance and scale. This book is a practical field guide to those less commonly used databases—vector, time-series, graph, columnar/OLAP, geospatial, event-stream, and more. It shows you how each engine works, where it excels, and how to choose the right engine for each workload. The book provides clear trade-off tables, indexing strategies, query patterns, and use-case recipes you can apply immediately. Starting with foundations in SQL and NoSQL, this book explores each “species” chapter by chapter, emphasizing hands-on scenarios and pitfalls to avoid. It then brings everything together in polyglot/hybrid architectures, helping you combine engines confidently, whether you're designing a similarity search with vector databases (ANN), building observability pipelines on time-series data and event logs, modeling



42,79 €
39,99 € (zzgl. MwSt.)

vorbestellbar, Erscheinungstermin ca.
Dezember 2026

Artikelnummer: 9798868827082

Medium: Buch

ISBN: 979-8-8688-2708-2

Verlag: APRESS L.P.

Erscheinungstermin: 02.09.2026

Sprache(n): Englisch

Auflage: 1. Auflage 2026

Produktform: Kartoniert

Seiten: 187

Format (B x H): 178 x 254 mm

