

Feifei Li · Xuan Zhou
Peng Cai · Rong Zhang
Gui Huang · XiangWen Liu

Cloud Native Database

Principle and Practice



Cloud Native Database

Feifei Li • Xuan Zhou • Peng Cai
Rong Zhang • Gui Huang • XiangWen Liu

Cloud Native Database

Principle and Practice



中国工信出版集团



电子工业出版社
PUBLISHING HOUSE OF ELECTRONICS INDUSTRY
<http://www.phei.com.cn>



Springer

Feifei Li
Cloud Intelligence Group
Alibaba Group (China)
Hangzhou, China

Xuan Zhou
School of Data Science & Engineering
East China Normal University
Shanghai, China

Peng Cai
School of Data Science & Engineering
East China Normal University
Shanghai, China

Rong Zhang
School of Data Science & Engineering
East China Normal University
Shanghai, China

Gui Huang
Cloud Intelligence Group
Alibaba Group (China)
Hangzhou, China

XiangWen Liu
Cloud Intelligence Group
Alibaba Group (China)
Hangzhou, China

ISBN 978-981-97-4056-7

ISBN 978-981-97-4057-4 (eBook)

<https://doi.org/10.1007/978-981-97-4057-4>

© Publishing House of Electronics Industry 2025

Jointly published with Publishing House of Electronics Industry

The print edition is not for sale in China (Mainland). Customers from China (Mainland) please order the print book from: Publishing House of Electronics Industry.

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publishers, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publishers nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publishers remain neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Singapore Pte Ltd.

The registered company address is: 152 Beach Road, #21-01/04 Gateway East, Singapore 189721, Singapore

If disposing of this product, please recycle the paper.

Cloud platform-based data management services have led to an increased need for cloud-native databases. This book, using PolarDB as an example, explains the principles and technologies of cloud-native databases from both theoretical and practical perspectives in a clear, concise, and easy-to-understand manner. I highly recommend this book as a great read.

Li Zhanhuai

Professor, Northwestern Polytechnical University; Director, CCF Technical
Committee on Databases

Cloud-native databases are the best choice for cloud computing platforms and have become a new favorite among users. This book, written by cloud computing and database experts, provides valuable content that can be used as a helpful reference.

Du Xiaoyong

Professor, Renmin University of China; Director, CCF Task Force on Big Data

Since its inception in the 1960s, databases have been recognized as a key infrastructure of the information society. With the development and popularization of the Internet over the past 20 years, the world has undergone profound changes. The future of the information society is becoming evident, and the digital transformation of the economy and society is ready to take off. The combination of databases and the Internet has brought forth the biggest challenge and opportunity for database development in the last decade. This has sparked renewed research and development efforts in the field. Cloud-native databases are the result of combining databases with cloud computing. They aim to provide database capabilities as services that are available to everyone, that is, to turn databases into public utilities. This is the first step in leveraging the power of data and providing a platform for

digital transformation. This book summarizes Alibaba Group's practical experiences in cloud-native databases. I believe it will help us gain a competitive edge in this emerging field.

Zhou Aoying
Professor and Vice President, East China Normal University;
Senior Database Scholar

Cloud computing platforms have made it possible for cloud-native databases to emerge and become widely used. This book explains the basic principles and core technologies of cloud-native databases from both theoretical and practical perspectives. The authors have extensive experience in researching database theory and have achieved significant results. This book has combined their research expertise with their experience in developing Alibaba Cloud's cloud-native database, PolarDB. I highly recommend it as a good read.

Peng Zhiyong
Professor and Associate Dean of the Big Data Institute, Wuhan University;
Deputy Director, CCF Technical Committee on Databases

Cloud-native databases have been a major innovation in the database field over the past decade, setting the trend for database development. This book explains the core technologies of cloud-native databases, such as computing-storage separation, log-as-data, and elastic multitenancy. It is an invaluable resource and deserves careful reading. This book is written by renowned experts from the academia and industry in the field of databases, and contains their insights into cloud-native databases.

Li Guoliang
Professor and Deputy Director of the Department of Computer Science and
Technology, Tsinghua University; Deputy Director, CCF Technical Committee on
Databases

Cloud-native databases are the most popular databases today, thanks to their excellent characteristics such as high scalability and availability. This book is a pioneering work on cloud-native databases, covering key theories and technical implementations. The authors are senior scholars and excellent professionals. I highly recommend this book to graduate students and developers who are interested in database technologies.

Cui Bin
Professor, Peking University; Deputy Director, CCF Technical Committee on
Databases

Data has become the key factor in the digital economy, and databases serve as the essential software infrastructure for storing and processing data. They play a crucial role in driving business development. Telecommunications service providers also highly value the development of database technologies. With the rapid development of cloud computing and big data, databases have evolved from traditional customized

deployments to on-demand, elastic, and scalable cloud services that are highly flexible and cost-effective. The book delves into Alibaba Cloud's technological exploration and practical experience in the cloudification of databases. I believe this book will offer valuable insights to readers, assisting them in successfully migrating to the cloud and expediting their digital transformation.

Chen Guo

Deputy General Manager, China Mobile Information Technology Co., Ltd

As technology advances and digital transformation progresses, data is gaining more attention as a valuable asset. Databases, as the foundation for storing and processing data, are also experiencing rapid development. Databases evolve in the following order: traditional databases, cloud-native databases, and generalized databases. Cloud-native databases are thriving with a brand-new technical architecture, which has contributed a lot to the implementation of cloud computing.

As the concept of cloud nativeness is still unclear to many, this book provides much-needed clarity on this topic. The authors possess profound theoretical knowledge, a visionary view of the industry, and extensive experience in the best practices of Alibaba's database products. Through their collaborative efforts, this book skillfully presents theoretical concepts and technical implementations. I highly recommend it to people who are interested in databases. I eagerly await the publication of this book.

Zhou Yanwei

Founder of Beijing Cloud-Ark Technology Co., Ltd. and Chief Architect of DTArk; Director of the Database Committee, China Computer Industry Association; Off-campus Supervisor, Zhejiang University

Databases have entered a new era of intense competition. As this book suggests, those who embrace cloud-nativeness are the ones who stand out. This book combines theory with practical insights, devoting many pages to technology selection. Drawing from the authors' extensive experience, it offers readers a comprehensive overview of cloud-native databases. It is undeniably a valuable read.

Zhang Wensheng

Chairman, PostgreSQL Chinese Community; Author of "A Practical Guide to PostgreSQL" and "PostgreSQL Guide: In-depth Exploration"

Foreword 1

Cloud-native databases are emerging as a new, vital form of databases. It was projected that 75% of databases will either be directly deployed in or migrated to the cloud by 2022. Alibaba Cloud's database products not only underpin the world's largest high-concurrency and low-latency e-commerce environment, providing seamless end-to-end online data management and services for millions of small-and medium-sized enterprises, but also offer stable and reliable data storage, processing, and analysis services for vital sectors, such as the government, manufacturing, finance, telecommunications, customs, transportation, and education industries. So far, Alibaba Cloud's database products have served over 100,000 enterprise users, enabling them to effortlessly access enterprise-level database services, substantially reduce costs, enhance operational efficiency, and generate novel business scenarios and value.

In response to the immense demands for concurrent data throughput and computational capabilities in e-commerce operations, Alibaba Cloud embarked on independent R&D of databases as early as 2010. We successfully tackled critical technological challenges such as storage-computing separation, distributed computing, high availability, compatibility, online-offline integration, and hybrid transactional/analytical processing (HTAP). We continued to upgrade the kernel and cluster architecture to meet the requirements of high-concurrency business scenarios, including the renowned Double 11 shopping festival. Moreover, we optimized our offerings for domestic chips and operating systems, laying a solid foundation for leveraging indigenous technology stacks and achieving complete autonomy and control. Along this journey, Alibaba Cloud's database solutions have garnered a range of prestigious accolades such as the award for the World's Leading Internet Scientific and Technological Achievements at the World Internet Conference, the first prize of Zhejiang Science and Technology Progress Award, and the first prize of the Science and Technology Progress Award of Chinese Institute of Electronics. Furthermore, Alibaba Cloud has become the first Chinese database vendor in Gartner's list of global database leaders. Standing at the threshold of the cloud computing era, Alibaba Cloud is committed to advancing

cloud-native database technologies in collaboration with our partners and developers. Together, we aim to create a sound industrial ecosystem to accelerate digital transformation across society.

DAMO Academy
Hangzhou, China

Jeff Zhang

Foreword II

Databases are among the most vital foundational systems in computer science. They serve as the bedrock software that supports the digital economy and occupy a pivotal strategic position. In the era of digital economy, effectively managing and leveraging diverse data resources is a prerequisite for scientific research and decision-making. The traditional database market is dominated by major commercial database vendors, forming a well-established database ecosystem. Databases store and process the core data resources of users, leading to high user stickiness and significant migration challenges. Due to a high degree of monopoly, Chinese database systems face fierce competition in the commercial market. The current national strategy places great emphasis on driving innovation and breakthroughs in fundamental technologies and explicitly prompts efforts to vigorously develop foundational software and advanced information technology services and expedite the development progress of database systems tailored for big data applications. Against this backdrop, database systems should not merely aim to replace existing products in the market, but to evolve and innovate, adapting to the emerging demands of cloud computing, big data, AI, and other new market trends. They should not only be useful but also user-friendly.

The rapid development of technologies like cloud computing has propelled foundational software toward a transformation journey into the cloud. An increasing number of enterprises are migrating their new applications to the cloud, intensifying the requirements for data storage and computational analysis capabilities. Cloud-native databases boast cloud elasticity, flexibility, and high availability, enabling them to provide robust innovation capabilities, diverse product offerings, economically efficient deployment methods, and pay-as-you-go payment models. Cloud-native distributed databases present a significant opportunity for groundbreaking innovation in the database domain. Cloudification opens up new possibilities for professionals working with databases.

The cloudification of databases has undergone two stages. The first stage is cloud hosting, which involves deploying existing database systems on cloud platforms to provide databases as on-demand services. The second stage is cloud-native implementation, where the hierarchical structure of databases is completely reconstructed

to leverage the resource pooling feature of the cloud. This decouples computing, storage, and networking resources to accommodate dynamic business needs by using elastic resource pools. In the second stage, databases are transformed comprehensively, unlocking opportunities for profound innovation. This book has emerged in response to this trajectory. As the primary author of this book, Dr. Li Feifei has dedicated over a decade to database research, followed by years of industry immersion focusing on the development of database systems. This book reflects a fusion of cutting-edge theory and practical engineering expertise. It provides a robust theoretical foundation and detailed technical implementation insights, to facilitate a deep understanding of the key technologies of cloud-native databases, including storage-computing separation, high availability, storage engines, distributed query engines, data distribution, and automatic load balancing. I believe this book will be invaluable to readers seeking to learn the latest database technologies.

Independent innovation in information infrastructure technologies is critical to informatization initiatives. However, this important task cannot be accomplished overnight. It necessitates a shared vision within society and close cooperation across the entire industry chain. Only by making unremitting efforts and seizing every opportunity presented by the evolving landscape we can succeed. I hope this book will inspire practitioners on their technological exploration journey and contribute to the development of new database systems.

Chinese Academy of Engineering
Beijing, China

Zuoning Chen

Foreword III

Database management systems are one of the most vital software systems in the field of computer science and technology. They serve as indispensable foundational platforms for information systems, providing essential support throughout the entire lifecycle of data from collection to classification, organization, processing, storage, analysis, and application. Without database management systems, informatization across industries would not be possible. Database management technologies have come a long way since their commercialization in the 1970s. Relational databases have established their dominant position in informatization, thanks to their concise conceptual framework, robust abstraction, powerful expressive capabilities, and transaction consistency guarantee, making them the “de facto standard” for data management technologies.

The widespread commercial use of the Internet has greatly accelerated the generation, circulation, and aggregation of data, raising new requirements and challenges for data management. The exponential growth in data volume necessitates better management practices, and the increasing diversity of data types calls for more flexible and diverse data models. The Internet, as the information infrastructure, places applications under greater pressure in terms of data throughput, concurrent access, and quick response to queries. The Internet is free of geographical constraints and therefore demands uninterrupted service provisioning from business systems. This imposes higher requirements on database systems in terms of availability and scalability, among other aspects. These new characteristics and scenarios have sparked a revolution in database technologies, giving rise to a multitude of innovative database technologies and products. The integration of distributed technologies has become a prominent feature, offering scalability and high availability to effectively address the demands of large-scale data processing and storage analysis.

The advent of cloud computing has ushered in a new phase for Internet applications. Computing and storage capabilities are offered to users as on-demand services. Besides infrastructure services, various underlying platform technologies, including databases and middleware, as well as application software, are also available as services. Cloud computing has triggered another wave of transformation in

database technologies. Backed by the application of distributed technologies, database services that boast real-time elastic scalability and geographically distributed availability are provided by leveraging the elastic resource pools of cloud platforms and decoupling the computation and storage layers of databases. This transformation is not only a technological evolution but also a renewal of business models, enabling users to access cost-effective, user-friendly, and highly available database services that are elastically scalable. This shift has given rise to the demand for cloud-native databases.

The reconstruction of database systems by using cloud-native technologies aligns with the technological trends and market demands. This book, primarily authored by Dr. Li Feifei, systematically explores the combination of cloud-native technologies and databases. It reviews the development of database management systems, outlining the major technological features of each important stage. It also traces the trajectory of database technologies, which culminate in the era of cloud-native databases. This book analyzes the trends in database technologies, deconstructs the technical stack of database systems, and explains the implementation of components such as shared storage systems, storage engines, and query engines after decoupling. It highlights how new technologies such as cloud-native architectures, distributed systems, high availability, and integration of software and hardware enhance the capabilities of database systems. This book also provides practical insights into the usage and O&M of cloud databases. With a well-structured and progressive layout, this book balances theory and practicality, making it a valuable reference for professionals in the field of database technologies.

The cloudification of databases is becoming a significant trend nowadays, presenting new opportunities for database management technologies and related industries. At the same time, it also calls for the cultivation of talents in this field. This book perfectly captures this zeitgeist.

Peking University
Beijing, China

Hong Mei

Preface

Background

For over six decades, database systems have been continuously developed to fulfill their role as one of the fundamental software components. As such, relational databases have dominated the market due to their strong data abstraction, expressive capabilities, and the easy-to-use SQL language. Over the past 50 years, the theory and technology of relational databases have come a long way. Numerous books have been published to delve into technical aspects such as SQL parsing, optimization and execution, transaction processing, log recovery, storage engines, and data dictionaries. Despite their maturity, database technologies continue to evolve due to a variety of factors, including the rapid development of the Internet and big data technology, complex business requirements, diverse data models, exponential data volumes, and advancements in hardware technologies.

Internet applications have completely reshaped people's lifestyles at an unprecedented pace, making an enormous amount of data available online. These data need to be stored, analyzed, and consumed, which, in turn, puts databases under unprecedented pressure. To adapt to the dynamic market, Internet applications quickly adjust their business forms and models. This leads to the emergence of more flexible and enriched data models and rapid changes in workload characteristics. To cope with such changes, databases must support elastic scaling to adapt to evolving business needs while keeping costs low. Traditional databases, often deployed as standalone systems with fixed specifications, struggle to meet these demands. This is where cloud computing comes in. By providing infrastructure as a service, cloud computing establishes large-scale resource pools and offers a unified virtualized abstraction interface. A massive operating system is established on diverse hardware by utilizing technologies such as containers, virtualization, orchestration, and microservices. Leveraging the capabilities of cloud computing, databases have transformed from fixed-specification instances to on-demand services, allowing users to access them as needed and scale them in real-time based on specific business requirements.

Cloud-native databases are not simply traditional databases deployed on cloud computing platforms. They have undergone a comprehensive transformation in terms of the overall architecture. They fully utilize the resource pooling capabilities of cloud computing platforms to decouple the previously monolithic databases, achieving complete separation of computing and storage resources. In addition, local storage is replaced by distributed cloud storage, and the computing layer becomes serverless. Cloud-native databases pool resources to support each layer of services, enabling independent and real-time scaling of resource pools to match dynamic workloads and maximize resource utilization.

Summary

This book portrays the evolution of database technologies in the era of cloud computing. Through specific examples, it illustrates how cloud-native and distributed technologies have enriched the essence of databases.

Chapter 1 offers a concise overview of database development. This chapter explains the structure, key modules, and implementation principles of typical relational databases. An SQL statement execution process is used as an example to illustrate these concepts.

Chapter 2 discusses the transformation of databases in the era of cloud computing, highlighting the evolution from standalone databases to cloud-native distributed databases. This chapter explores the technical changes brought by cloud computing and examines the potential trends in database technologies.

Chapter 3 focuses on the architectural design principles of cloud-native databases and the reasons behind these principles. Additionally, this chapter analyzes the technical features of several prominent cloud-native databases in the market, such as AWS Aurora, Alibaba Cloud PolarDB, and Microsoft Socrates.

Chapters 4–7, respectively, delve into the implementation principles of important components of cloud-native databases, including storage engines, shared storage, database caches, and computing engines. Each chapter follows the same structure, in which the theoretical foundations and general implementation methods of the components are explained, and then targeted improvements and optimization methods specific to cloud-native databases are introduced.

Chapter 8 provides a detailed explanation of distributed database technologies that support scale-in and scale-out, including their application and implementation principles. This chapter also highlights how the integration of database technologies with cloud-native technologies takes the database technologies to new heights.

Chapters 9 and 10 center around the practical applications of cloud-native databases. By using PolarDB as an example, these chapters cover relevant topics, such as creating database instances in the cloud, optimizing usage and O&M, and harnessing the elastic, high availability, security, and cost-effectiveness features offered by cloud databases.

Primary Authors

This book is authored by Dr. Li Feifei from Alibaba Cloud and Professor Zhou Xuan from East China Normal University (ECNU). Some of the content is contributed by Professor Cai Peng and Professor Zhang Rong from ECNU and also senior technical expert Huang Gui from Dr. Li Feifei's team. Liu Xiangwen, the Vice President of Alibaba Cloud and General Manager for marketing, Alibaba Cloud Intelligence has also made significant contributions. Other technical experts from Alibaba Cloud's database team, including Zhang Yingqiang, Wang Jianying, Hu Qingda, Chen Zongzhi, Wang Yuhui, Wang Bo, Sun Yue, Zhuang Zechao, Ying Shanshan, Song Zhao, Wang Kang, Cheng Xuntao, Zhang Haiping, Wu Xiaofei, Wu Xueqiang, Yang Shukun, and others, have provided valuable technical materials, and we sincerely appreciate their contributions.

Special thanks to Jeff Zhang, Managing Director of DAMO Academy, Academician Chen Zuoning from the Chinese Academy of Engineering, and Academician Mei Hong from the Chinese Academy of Sciences for writing the forewords for this book.

We would also like to express our gratitude to Professor Li Zhanhuai, Professor Du Xiaoyong, Professor Zhou Aoying, Professor Peng Zhiyong, Professor Li Guoliang, Professor Cui Bin, General Manager Chen Guo, President Zhou Yanwei, and Chairman Zhang Wensheng for their testimonials.

Additionally, we extend our appreciation to the technical experts from Alibaba Cloud's database team, including Huang Gui, Yang Xinjun, Lou Jiangang, You Tianyu, Wu Wenqi, Chen Zongzhi, Liang Chen, Zhang Yingqiang, Wang Jianying, Hu Qingda, Weng Ninglong, Fu Dachao, Fu Cuiyun, Wang Yuhui, Yuan Lixiang, Sun Jingyuan, Cai Chang, Zhou Jie, Xu Jiawei, Wu Xiaofei, Xie Rongbiao, Wang Kang, Zheng Song, Ren Zhuo, Wei Zetao, Sun Yuxuan, Zhang Xing, Li Ziqiang, Xu Dading, Xiong Meihui, Liang Gaozhong, Chen Shiyang, Chen Jiang, Xu Jie, Cai Xin, Yu Nanlong, Wang Yujie, Chen Shibin, Wu Qianqian, Sun Yue, Zhao Minghuan, Sun Haiqing, Li Wei, Yang Yuming, and Han Ken for their contributions to the translation of textbooks. Thanks to Wang Yuan and Xiao Simiao from the Alibaba Cloud's database team for their contributions in organizing the translation of the textbook.

This book would not have been possible without the collective efforts of everyone involved.

As it was completed within a limited timeframe, this book may not answer everything there is to know about database systems. With this, we encourage readers to kindly share their feedback.

Hangzhou, China
 Shanghai, China
 Shanghai, China
 Shanghai, China
 Hangzhou, China
 Hangzhou, China
 February 2024

Feifei Li
 Xuan Zhou
 Peng Cai
 Rong Zhang
 Gui Huang
 XiangWen Liu

Introduction

This book thoroughly analyzes the technological evolution of databases, which serve as core software systems, in the era of cloud computing. It explores the progressive development of traditional database technologies toward cloud-native architectures from various perspectives, including architectural design, implementation mechanisms, and system optimization. On the basis of the fusion of theory and practice, this book examines SQL optimization and execution, transaction processing, caching, indexing, and other features employed by widely used database systems like MySQL and PostgreSQL. It also explores the trade-offs and compromises made to meet practical application requirements, the improvements made to suit complex scenarios, and the underlying rationale behind these choices. Furthermore, this book draws on Alibaba Cloud’s database research and development experience, highlighting the core technical principles that have enabled modern databases to evolve into services, such as cloud computing resource pooling and distributed technologies for high availability, elastic scaling, and on-demand usage.

This book offers comprehensive theoretical knowledge and practical experience by navigating the latest trends in database development, thereby inspiring readers to delve deeper into the subject. It can serve as a textbook for undergraduate and graduate students majoring in information-related disciplines at higher educational institutions, as well as a reference book for professionals engaged in kernel development and system O&M in the database industry.

Contents

1	Database Development Milestones	1
1.1	Overview of Database Development	1
1.1.1	Emergence	2
1.1.2	Commercialization	2
1.1.3	Maturation	3
1.1.4	Cloud-Native and Distributed Era	3
1.2	Database Technology Development Trends	5
1.2.1	Cloud-Native and Distributed Architectures	5
1.2.2	Integration of Big Data and Databases	6
1.2.3	Hardware-Software Integration	7
1.2.4	Multimodality	7
1.2.5	Intelligent O&M	8
1.2.6	Security and Trust	8
1.3	Key Components of Relational Databases	9
1.3.1	Access Management Component	9
1.3.2	Query Engine	10
1.3.3	Transaction Processing System	15
1.3.4	Storage Engine	18
	References	21
2	Database and Cloud Nativeness	23
2.1	Development of Databases in the Cloud Era	23
2.1.1	Rise of Cloud Computing	23
2.1.2	Database as a Service	24
2.2	Challenges Faced by Databases in the Cloud-Native Era	26
2.3	Characteristics of Cloud-Native Databases	27
2.3.1	Layered Architecture	27
2.3.2	Resource Decoupling and Pooling	27
2.3.3	Elastic Scalability	27
2.3.4	High Availability and Data Consistency	28

- 2.3.5 Multitenancy and Resource Isolation. 29
- 2.3.6 Intelligent O&M. 29
- References. 30
- 3 Architecture of Cloud-Native Database 31**
 - 3.1 Design Principles 31
 - 3.1.1 Essence of Cloud-Native Databases 31
 - 3.1.2 Separation of Computing and Storage 32
 - 3.2 Architecture Design 33
 - 3.3 Typical Cloud-Native Databases 35
 - 3.3.1 AWS Aurora. 35
 - 3.3.2 PolarDB 41
 - 3.3.3 Microsoft Socrates 45
 - References. 49
- 4 Storage Engine 51**
 - 4.1 Data Organization 51
 - 4.1.1 B+ Tree 52
 - 4.1.2 B+ Tree in the InnoDB Engine 54
 - 4.1.3 LSM-Tree. 58
 - 4.2 Concurrency Control 62
 - 4.2.1 Basic Concepts. 62
 - 4.2.2 Lock-Based Concurrency Control 62
 - 4.2.3 Timestamp-Based Concurrency Control 64
 - 4.2.4 MVCC 67
 - 4.2.5 Implementation of MVCC in InnoDB. 69
 - 4.3 Logging and Recovery 72
 - 4.3.1 Basic Concepts. 72
 - 4.3.2 Logical Logs 72
 - 4.3.3 Physical Logs. 73
 - 4.3.4 Recovery Principles 74
 - 4.3.5 Binlog of MySQL 74
 - 4.3.6 Physical Logs of InnoDB. 75
 - 4.4 LSM-Tree Storage Engine 77
 - 4.4.1 PolarDB X-Engine. 77
 - 4.4.2 High-Performance Transaction Processing 79
 - 4.4.3 Hardware-Facilitated Software Optimization 82
 - 4.4.4 Cost-Effective Tiered Storage 86
 - 4.4.5 Dual Storage Engine Technology 92
 - 4.4.6 Experimental Evaluation 93
 - References. 97
- 5 High-Availability Shared Storage System 99**
 - 5.1 Basics of High Availability 99
 - 5.1.1 Leader and Follower Replicas 100
 - 5.1.2 Quorum 101

- 5.1.3 Paxos 102
- 5.1.4 Raft 105
- 5.1.5 Parallel Raft 108
- 5.2 High Availability of Clusters 110
 - 5.2.1 High Availability of MySQL Clusters 110
 - 5.2.2 High Availability of PolarDB 114
- 5.3 Shared Storage Architectures 129
 - 5.3.1 Aurora Storage System 130
 - 5.3.2 PolarFS 131
- 5.4 File System Optimization 134
 - 5.4.1 User Space I/O Computing 134
 - 5.4.2 Near-Storage Computing 137
- References 143
- 6 Database Cache 145**
 - 6.1 Introduction to the Database Cache 145
 - 6.1.1 Role of the Database Cache 145
 - 6.1.2 Buffer Pool 146
 - 6.2 Cache Recovery 146
 - 6.2.1 Challenges of Caching in the Cloud Environment 146
 - 6.2.2 Cache Recovery Based on CPU-Memory Separation 147
 - 6.3 PolarDB Practices 149
 - 6.3.1 Optimization of the Buffer Pool 149
 - 6.3.2 Optimization of the Data Dictionary Cache and the File System Cache 154
 - 6.3.3 RDMA-Based Shared Memory Pool 155
 - References 160
- 7 Computing Engine 163**
 - 7.1 Overview of Query Processing 163
 - 7.1.1 Overview of Database Query Processing 163
 - 7.1.2 Overview of Parallel Queries 165
 - 7.2 Query Execution Models 167
 - 7.2.1 Volcano Model 168
 - 7.2.2 Compiled Execution Model 168
 - 7.2.3 Vectorized Execution Model 169
 - 7.3 Overview of Query Optimization 169
 - 7.3.1 Introduction to Query Optimization 169
 - 7.3.2 Logical Query Optimization 170
 - 7.3.3 Physical Query Optimization 170
 - 7.3.4 Other Optimization Methods 171
 - 7.4 Practical Application of PolarDB Query Engine 171
 - 7.4.1 Parallel Query Technology in PolarDB 172
 - 7.4.2 Execution Plan Management in PolarDB 187
 - 7.4.3 Vectorized Execution in PolarDB 194
 - References 197