

# Expert Techniques To Build Scalable Reliable And Fault Tolerant Database

Building a scalable, reliable, and fault-tolerant database is crucial for businesses in today's data-driven world. With increasing amounts of data being generated each day, it becomes essential to have a database that can handle high volumes of traffic, provide efficient storage, and ensure data integrity. In this article, we will explore some expert techniques that can help you achieve these goals.

## 1. Data Sharding

Data sharding is a technique that involves dividing a large database into smaller, more manageable parts called shards. Each shard contains a subset of the data and can be stored on separate servers. By distributing the data across multiple servers, you can increase the overall capacity of your database and improve its performance. This technique is particularly useful for handling large datasets and high traffic loads.

When implementing data sharding, it's essential to carefully choose the shard key, which determines how the data will be distributed among the shards. The shard key should be selected based on the access patterns of your application to ensure an even distribution of data and efficient query execution.



## Mastering PostgreSQL 11: Expert techniques to build scalable, reliable, and fault-tolerant database applications, 2nd Edition

by Hans-Jürgen Schönig (2nd Edition, Kindle Edition)

★★★★☆ 4.2 out of 5

Language : English

File size : 1206 KB

Text-to-Speech : Enabled

Screen Reader : Supported  
Enhanced typesetting : Enabled  
Print length : 448 pages



## 2. Replication

Database replication involves creating multiple copies of your database on separate servers. This technique provides redundancy and ensures that if one server fails, the other servers can continue serving data. Replication can also help improve read performance by allowing read operations to be performed on any replica server.

There are different types of replication, such as master-slave replication and master-master replication. In master-slave replication, one server (the master) handles write operations, while the other servers (the slaves) replicate the changes. In master-master replication, multiple servers can handle both read and write operations and synchronize with each other. The choice of replication method depends on your specific requirements and the level of fault tolerance you need.

## 3. Load Balancing

Load balancing is the process of distributing incoming network traffic across multiple servers to ensure optimal resource utilization and prevent any single server from becoming overwhelmed. Load balancers act as intermediaries between clients and servers, directing requests to the server with the least load.

Implementing load balancing for your database can help distribute the workload evenly among the servers, ensuring scalability and fault tolerance. It also

improves performance by preventing any single server from being overloaded. There are various load balancing algorithms available, such as round-robin, least connections, and weighted round-robin, which can be used depending on your specific requirements.

#### **4. Automatic Failover**

In a distributed database environment, failures are inevitable. To ensure high availability and minimal downtime, it's crucial to implement automatic failover mechanisms. Automatic failover allows the system to seamlessly switch to a backup or replica server in case of a primary server failure.

When configuring automatic failover, it's important to set up monitoring and alerting systems that can detect server failures and trigger the failover process. This can be achieved through various techniques such as heartbeat monitoring, where a heartbeat signal is sent between servers, and if the signal is not received within a specified time frame, a failover is initiated.

#### **5. Consistent Hashing**

Consistent hashing is a technique used to distribute data evenly across multiple servers while minimizing the impact of adding or removing servers from the system. With consistent hashing, the hash space is divided into intervals, and each server is responsible for a range of hash values.

Consistent hashing ensures that when a server is added or removed, only a small portion of the data needs to be remapped, minimizing the amount of data migration required. This technique is particularly useful in dynamically scalable environments where servers need to be added or removed frequently.

#### **6. Data Compression and Caching**

Data compression and caching are techniques that can help optimize database performance and reduce storage requirements. Data compression reduces the size of stored data, resulting in reduced disk I/O and improved query performance. Caching, on the other hand, stores frequently accessed data in memory, reducing the need to retrieve data from disk.

By implementing effective data compression and caching strategies, you can reduce the overall storage costs and improve the response time of your database. However, it's important to strike the right balance between compression ratios and query performance, as excessive compression may impact query execution time.

## **7. Regular Database Maintenance**

To ensure the long-term reliability and performance of your database, regular maintenance is essential. This includes tasks such as monitoring database health, optimizing database queries, performing backups, and applying software patches and updates.

By regularly monitoring and optimizing your database, you can identify and fix any performance bottlenecks, improve query execution time, and ensure data integrity. Additionally, having a robust backup and recovery strategy in place can protect your data from accidental deletion, system failures, or security breaches.

Building a scalable, reliable, and fault-tolerant database requires implementing a combination of techniques such as data sharding, replication, load balancing, automatic failover, consistent hashing, data compression, caching, and regular maintenance. By carefully considering your specific requirements and leveraging these techniques, you can ensure that your database can handle high traffic loads, provide efficient storage, and maintain data integrity.

Remember, the success of your business often relies on the reliability and performance of your database, so investing time and effort into building a robust database infrastructure is crucial. Implementing these expert techniques will help you build a database that can scale with your business, withstand failures, and deliver optimal performance.



## Mastering PostgreSQL 11: Expert techniques to build scalable, reliable, and fault-tolerant database applications, 2nd Edition

by Hans-Jürgen Schönig (2nd Edition, Kindle Edition)

★★★★☆ 4.2 out of 5

Language : English

File size : 1206 KB

Text-to-Speech : Enabled

Screen Reader : Supported

Enhanced typesetting : Enabled

Print length : 448 pages



Master the capabilities of PostgreSQL 11 to efficiently manage and maintain your database

### Key Features

- Master advanced concepts of PostgreSQL 11 with real-world datasets and examples
- Explore query parallelism, data replication, and database performance while working with larger datasets
- Extend the functionalities of your PostgreSQL instance to suit your organization's needs with minimal effort

## **Book Description**

This second edition of Mastering PostgreSQL 11 helps you build dynamic database solutions for enterprise applications using the latest release of PostgreSQL, which enables database analysts to design both the physical and technical aspects of the system architecture with ease.

This book begins with an to the newly released features in PostgreSQL 11 to help you build efficient and fault-tolerant PostgreSQL applications. You'll examine all of the advanced aspects of PostgreSQL in detail, including logical replication, database clusters, performance tuning, monitoring, and user management. You will also work with the PostgreSQL optimizer, configuring PostgreSQL for high speed, and see how to move from Oracle to PostgreSQL. As you progress through the chapters, you will cover transactions, locking, indexes, and optimizing queries to improve performance. Additionally, you'll learn to manage network security and explore backups and replications, while understanding the useful extensions of PostgreSQL so that you can optimize the speed and performance of large databases.

By the end of this book, you will be able to use your database to its utmost capacity by implementing advanced administrative tasks with ease.

## **What you will learn**

- Get to grips with advanced PostgreSQL 11 features and SQL functions
- Make use of the indexing features in PostgreSQL and fine-tune the performance of your queries
- Work with stored procedures and manage backup and recovery
- Master replication and failover techniques

- Troubleshoot your PostgreSQL instance for solutions to common and not-so-common problems
- Perform database migration from MySQL and Oracle to PostgreSQL with ease

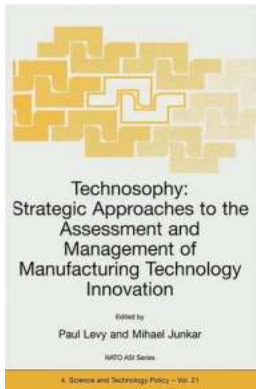
## **Who this book is for**

This book is for data and database professionals wanting to implement advanced functionalities and master complex administrative tasks with PostgreSQL 11. Prior experience of database administration with PostgreSQL database will aid in understanding the concepts covered in this book.

## **Table of Contents**

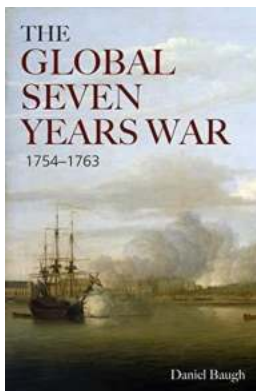
1. PostgreSQL Overview
2. Understanding Transactions and Locking
3. Making Use of Indexes
4. Handling Advanced SQL
5. Log Files and System Statistics
6. Optimizing Queries for Good Performance
7. Writing Stored Procedures
8. Managing PostgreSQL Security
9. Handling Backup and Recovery
10. Making Sense of Backups and Replication
11. Deciding on Useful Extensions
12. Troubleshooting PostgreSQL

## 13. Migrating to PostgreSQL



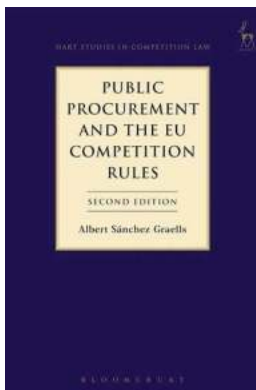
### **The Ultimate Guide to Strategic Approaches in Manufacturing Management: Boosting Efficiency and Profitability**

In the ever-evolving manufacturing industry, staying ahead of the competition requires strategic approaches to both assessment and management. The key to success lies in...



### **The Explosive Global Conflict: The Seven Years War (1754-1763)**

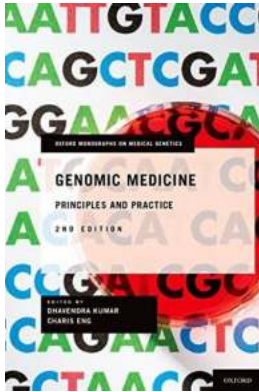
The Global Seven Years War, lasting from 1754 to 1763, was a fierce and expansive conflict that reshaped the world order. Fought on multiple continents...



### **Public Procurement And The Eu Competition Rules: An In-Depth Analysis**

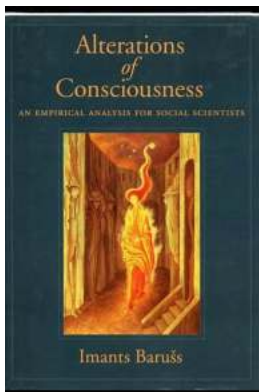
Public procurement is a critical aspect of the economic activities of any country. It involves the acquisition of goods, services, and construction works by public...





## Unlocking the Secrets: Principles And Practice Oxford Monographs On Medical Genetics 65

In the world of medical genetics, staying informed about the latest research and practices is essential to providing the best care for patients. One invaluable resource that...



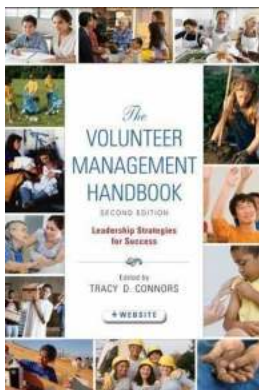
## An Empirical Analysis For Social Scientists: Unveiling the Patterns of Human Behavior

As social scientists, our aim is to understand and explain human behavior. With the advancements in technology and the increasing availability of data, we now have the tools...



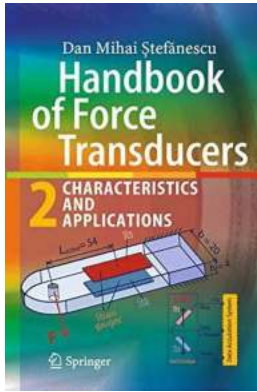
## Expert Techniques To Build Scalable Reliable And Fault Tolerant Database

Building a scalable, reliable, and fault-tolerant database is crucial for businesses in today's data-driven world. With increasing amounts of data being generated...



## Unlocking the Secrets of Leadership Strategies for Success in the Nonprofit Sector

In today's fast-paced and ever-changing world, effective leadership is vital for the success of any organization, especially in the nonprofit sector. Wiley Nonprofit Law...



# The Ultimate Handbook of Force Transducers: Unveiling the Characteristics and Applications

Unlocking the Secrets of Force Measurement for Precise Engineering

Force is a fundamental aspect of physics that requires precise...