

## High Availability With Postgresql And Pacemaker

Master the capabilities of PostgreSQL 9.6 to efficiently manage and maintain your database

**About This Book\*** Your one-stop guide to mastering the advanced concepts in PostgreSQL with ease\* Master query optimization, replication, and high availability with PostgreSQL\* Extend the functionalities of PostgreSQL to suit your organizational needs with minimum effort

**Who This Book Is For** If you are a PostgreSQL data architect or an administrator who wants to understand how to implement advanced functionalities and master complex administrative tasks with PostgreSQL, then this book is perfect for you. Prior experience of administrating a PostgreSQL database and a working knowledge of SQL is required to make the best use of this book.

**What You Will Learn\*** Get to grips with the advanced features of PostgreSQL 9.6 and handle advanced SQL\* Make use of the indexing features in PostgreSQL and fine-tune the performance of your queries\* Work with the stored procedures and manage backup and recovery\* Master the replication and failover techniques\* Troubleshoot your PostgreSQL instance for solutions to the common and not-so-common problems\* Learn how to migrate your database from MySQL and Oracle to PostgreSQL without any hassle

**In Detail** PostgreSQL is an open source database used for handling large datasets (Big Data) and as a JSON document database. It also has applications in the software and web domains. This book will enable you to build better PostgreSQL applications and administer databases more efficiently. We begin by explaining the advanced database design concepts in PostgreSQL 9.6, along with indexing and query optimization. You will also see how to work with event triggers and perform concurrent transactions and table partitioning, along with exploring SQL and server tuning. We will walk you through implementing advanced administrative tasks such as server maintenance and monitoring, replication, recovery and high availability, and much more. You will understand the common and not-so-common troubleshooting problems and how you can overcome them. By the end of this book, you will have an expert-level command of the advanced database functionalities and will be able to implement advanced administrative tasks with PostgreSQL.

**Style and Approach** This book is a comprehensive guide covering all the concepts you need to master PostgreSQL. Packed with hands-on examples, tips and tricks, even the most advanced concepts are explained in a very easy-to-follow manner. Every chapter in the book does not only focus on how each task is performed, but also why.

Develop, deploy, and scale your applications with Google Cloud Platform

**Key Features** Create and deploy your applications on Google Cloud Platform Store and manage source code and debug Cloud-hosted apps with plugins and IDEs Streamline developer workflows with tools for alerting and managing deployments

**Book Description** Google Cloud Platform (GCP) provides autoscaling compute power and distributed in-memory cache, task queues, and datastores to

write, build, and deploy Cloud-hosted applications. With Google Cloud Platform for Developers, you will be able to develop and deploy scalable applications from scratch and make them globally available in almost any language. This book will guide you in designing, deploying, and managing applications running on Google Cloud. You'll start with App Engine and move on to work with Container Engine, compute engine, and cloud functions. You'll learn how to integrate your new applications with the various data solutions on GCP, including Cloud SQL, Bigtable, and Cloud Storage. This book will teach you how to streamline your workflow with tools such as Source Repositories, Container Builder, and StackDriver. Along the way, you'll see how to deploy and debug services with IntelliJ, implement continuous delivery pipelines, and configure robust monitoring and alerting for your production systems. By the end of this book, you'll be well-versed with all the development tools of Google Cloud Platform, and you'll develop, deploy, and manage highly scalable and reliable applications. What you will learn Understand the various service offerings on GCP Deploy and run services on managed platforms such as App Engine and Container Engine Securely maintain application states with Cloud Storage, Datastore, and Bigtable Leverage StackDriver monitoring and debugging to minimize downtime and mitigate issues without impacting users Design and implement complex software solutions utilizing Google Cloud Integrate with best-in-class big data solutions such as Bigquery, Dataflow, and Pub/Sub Who this book is for Google Cloud Platform for Developers is for application developers. This book will enable you to fully leverage the power of Google Cloud Platform to build resilient and intelligent software solutions.

"Introduction. 1. Pt. I. Getting Started. 3. 1. Getting Started. 5. 2. Setting Up PHP. 15. 3. PHP Basics. 43. 4. Object-Oriented PHP. 121. Pt. II. Getting Started with PostgreSQL. 137. 5. Relational and Object-Relational Database Concepts. 139. 6. Installing PostgreSQL. 149. 7. Basic SQL. 177. 8. Advanced SQL. 225. 9. Embedded Languages. 307. 10. PostgreSQL Administration. 349. Pt. III. PHP/PostgreSQL Interaction. 379. 11. Writing Database-Driven Applications. 381. 12. Working with BLOBs. 405. 13. Working with Persistent Database Connections. 423. Pt. IV. Advanced Technologies. 433. 14. Managing Regular Expressions. 435. 15. Session Management. 463. 16. Working with Dynamic Documents, Images, and Movies. 475. 17. Working with Dates and Time. 521. 18. Tuning. 551. 19. XML. 573. 20. Security Issues. 585. Pt. V. Practical Examples. 597. 21. Web Applications. 599. 22. Extending PostgreSQL. 669. 23. High-Availability Systems. 695. Pt. VI. Migration. 709. 24. Migration. 711. . Index. 721.

Welcome to the "PostgreSQL 8.4 Official Documentation - Volume II. Server Administration"! After many years of development, PostgreSQL has become feature-complete in many areas. This release shows a targeted approach to adding features (e.g., authentication, monitoring, space reuse), and adds capabilities defined in the later SQL standards. "PostgreSQL" leads users through the internals of an open-source database. Throughout the book are explanations of

data structures and algorithms, each backed by a concrete example from the actual source code. Each section contains information about performance implications, debugging techniques, and pointers to more information (on the Web and in book form).

Write optimized queries. This book helps you write queries that perform fast and deliver results on time. You will learn that query optimization is not a dark art practiced by a small, secretive cabal of sorcerers. Any motivated professional can learn to write efficient queries from the get-go and capably optimize existing queries. You will learn to look at the process of writing a query from the database engine's point of view, and know how to think like the database optimizer. The book begins with a discussion of what a performant system is and progresses to measuring performance and setting performance goals. It introduces different classes of queries and optimization techniques suitable to each, such as the use of indexes and specific join algorithms. You will learn to read and understand query execution plans along with techniques for influencing those plans for better performance. The book also covers advanced topics such as the use of functions and procedures, dynamic SQL, and generated queries. All of these techniques are then used together to produce performant applications, avoiding the pitfalls of object-relational mappers.

**What You Will Learn**

- Identify optimization goals in OLTP and OLAP systems
- Read and understand PostgreSQL execution plans
- Distinguish between short queries and long queries
- Choose the right optimization technique for each query type
- Identify indexes that will improve query performance
- Optimize full table scans
- Avoid the pitfalls of object-relational mapping systems
- Optimize the entire application rather than just database queries

**Who This Book Is For**

IT professionals working in PostgreSQL who want to develop performant and scalable applications, anyone whose job title contains the words "database developer" or "database administrator" or who is a backend developer charged with programming database calls, and system architects involved in the overall design of application systems running against a PostgreSQL database

**Master PostgreSQL 12 features** such as advanced indexing, high availability, monitoring, and much more to efficiently manage and maintain your database

**Key Features**

- Grasp advanced PostgreSQL 12 concepts with real-world examples and sample datasets
- Explore query parallelism, data replication, database administration, and more
- Extend PostgreSQL functionalities to suit your organization's needs with minimal effort

**Book Description**

Thanks to its reliability, robustness, and high performance, PostgreSQL has become the most advanced open source database on the market. This third edition of *Mastering PostgreSQL* helps you build dynamic database solutions for enterprise applications using the latest release of PostgreSQL, which enables database analysts to design both physical and technical aspects of system architecture with ease. Starting with an introduction to the newly released features in PostgreSQL 12, this book will help you build efficient and fault-tolerant PostgreSQL applications. You'll thoroughly examine the advanced features of

PostgreSQL, including logical replication, database clusters, performance tuning, monitoring, and user management. You'll also work with the PostgreSQL optimizer, configure PostgreSQL for high speed, and understand how to move from Oracle to PostgreSQL. As you progress through the chapters, you'll cover transactions, locking, indexes, and how to optimize queries for improved performance. Additionally, you'll learn how to manage network security and explore backups and replications while understanding useful PostgreSQL extensions to help you in optimizing the performance of large databases. By the end of this PostgreSQL book, you'll be able to get the most out of your database by implementing advanced administrative tasks effortlessly. What you will learn Understand the advanced SQL functions in PostgreSQL 12 Use indexing features in PostgreSQL to fine-tune the performance of queries Work with stored procedures and manage backup and recovery Master replication and failover techniques to reduce data loss Replicate PostgreSQL database systems to create backups and to scale your database Manage and improve the security of your server to protect your data Troubleshoot your PostgreSQL instance for solutions to common and not-so-common problems Who this book is for This book is for PostgreSQL developers and administrators and database professionals who want to implement advanced functionalities and master complex administrative tasks with PostgreSQL 12. Prior exposure to PostgreSQL as well as familiarity with the basics of database administration is expected.

In Detail PostgreSQL is a powerful, open source, enterprise database. This PostgreSQL 9 Administration Cookbook LITE book describes key aspects of the PostgreSQL open source database system. The book will help a sysadmin or DBA with key administration issues in PostgreSQL: configuration, monitoring and diagnosis, and setting up regular maintenance. This hands-on guide will assist developers working on live databases supporting web or enterprise software applications. To find out more about upgrading to the full edition, visit [www.packtpub.com/lite-editions](http://www.packtpub.com/lite-editions) and log into your account for offers and help. If you don't have an account on PacktPub.com, visit today and set one up! Approach Written in the cookbook style, this book offers learning and techniques through recipes. It contains step-by-step instructions for administrators and developers to manage databases in PostgreSQL. The book is designed in such a way that you can read it chapter by chapter or refer to recipes in no particular order. Who this book is for This book is for Sysadmins, Database Administrators, Architects, Developers, and anyone with an interest in planning for or running live production databases. The book assumes that you are familiar with the basic operation of PostgreSQL.

Enterprises require support and agility to work with big data repositories and relational databases. FUJITSU Enterprise Postgres is one of the leading relational database management systems (RDBMSs), and it is designed to work with large data sets. As more companies transform their infrastructures with hybrid cloud services, they require environments that protect the safety of their data and business rules. At IBM®, we believe that your data is yours and yours alone. The

insights and advantages that come from your data are yours to use in the pursuit of your business objectives. IBM is dedicated to this mission, and the IBM LinuxONE platform is designed around this core statement. IBM LinuxONE is a secure and scalable data serving and computing platform that is made for today's critical workloads. IBM LinuxONE is an all-Linux enterprise platform for open innovation that combines the best of Linux and open technology with the best of enterprise computing in one system. Combining FUJITSU Enterprise Postgres, which is a robust Relational Database Management System (RDBMS) that provides strong query performance and high availability (HA), with IBM LinuxONE can transform your application and data portfolio by providing innovative data privacy, security, and cyber resiliency capabilities, which are all delivered with minimal downtime. This IBM Redbooks® publication describes data serving with FUJITSU Enterprise Postgres 12 that is deployed on IBM LinuxONE, which provides the scalability, business-critical availability, and security that your enterprise requires. This publication is useful to IT architects, system administrators, and others who are interested in understanding the significance of using FUJITSU Enterprise Postgres on IBM LinuxONE. This publication is written for those who are familiar with IBM LinuxONE and have some experience in the use of PostgreSQL.

\*The most updated PostgreSQL book on the market, covering version 8.0 \*Highlights the most popular PostgreSQL APIs, including C, Perl, PHP, and Java \*This is two books in one; it simultaneously covers key relational database design principles, while teaching PostgreSQL

Build cost-effective and robust cloud solutions with Google Cloud Platform (GCP) using these simple and practical recipes Key Features Explore the various service offerings of the GCP Host a Python application on Google Compute Engine Securely maintain application states with Cloud Storage, Datastore, and Bigtable Book Description GCP is a cloud computing platform with a wide range of products and services that enable you to build and deploy cloud-hosted applications. This Learning Path will guide you in using GCP and designing, deploying, and managing applications on Google Cloud. You will get started by learning how to use App Engine to access Google's scalable hosting and build software that runs on this framework. With the help of Google Compute Engine, you'll be able to host your workload on virtual machine instances. The later chapters will help you to explore ways to implement authentication and security, Cloud APIs, and command-line and deployment management. As you hone your skills, you'll understand how to integrate your new applications with various data solutions on GCP, including Cloud SQL, Bigtable, and Cloud Storage. Following this, the book will teach you how to streamline your workflow with tools, including Source Repositories, Container Builder, and Stackdriver. You'll also understand how to deploy and debug services with IntelliJ, implement continuous delivery pipelines, and configure robust monitoring and alerts for your production systems. By the end of this

Learning Path, you'll be well versed with GCP's development tools and be able to develop, deploy, and manage highly scalable and reliable applications. This Learning Path includes content from the following Packt products: Google Cloud Platform for Developers Ted Hunter and Steven Porter Google Cloud Platform Cookbook by Legorie Rajan PS What you will learn Host an application using Google Cloud Functions Migrate a MySQL database to Cloud Spanner Configure a network for a highly available application on GCP Learn simple image processing using Storage and Cloud Functions Automate security checks using Policy Scanner Deploy and run services on App Engine and Container Engine Minimize downtime and mitigate issues with Stackdriver Monitoring and Debugger Integrate with big data solutions, including BigQuery, Dataflow, and Pub/Sub Who this book is for This Learning Path is for IT professionals, engineers, and developers who want to implement Google Cloud in their organizations. Administrators and architects planning to make their organization more efficient with Google Cloud will also find this Learning Path useful. Basic understanding of GCP and its services is a must.

PostgreSQL has become the most advanced open source database on the market. This book adopts a step-by-step approach to meet almost every requirement you can think of while deploying PostgreSQL in production environments. You will not only learn how to design and manage your database but also discover how to administer and secure the database.

This book will get you up and running with the working of relational databases, data modeling, data manipulation, and more. You will learn to build efficient relational database solutions from scratch using the latest features of PostgreSQL 12 and 13. You'll also be able to identify bottlenecks to enhance the performance of database applications.

Get to know effective ways to improve PostgreSQL's performance and master query optimization, and database monitoring. About This Book Perform essential database tasks such as benchmarking the database and optimizing the server's memory usage Learn ways to improve query performance and optimize the PostgreSQL server Explore a wide range of high availability and replication mechanisms to build robust, highly available, scalable, and fault-tolerant PostgreSQL databases Who This Book Is For If you are a developer or administrator with limited PostgreSQL knowledge and want to develop your skills with this great open source database, then this book is ideal for you. Learning how to enhance the database performance is always an exciting topic to everyone, and this book will show you enough ways to enhance the database performance. What You Will Learn Build replication strategies for homogeneous and heterogeneous databases Test and build a powerful machine with multiple bench marking techniques Get to know a few SQL injection techniques Find out how to manage the replication using multiple tools Benchmark the database server using multiple strategies Work with the query processing algorithms and their internal behaviors Build a proper plan to

upgrade or migrate to PostgreSQL from other databases See the essential database load balancing techniques and the various partitioning approaches PostgreSQL provides Learn memory optimization techniques and database server configurations In Detail PostgreSQL is one of the most powerful and easy to use database management systems. It has strong support from the community and is being actively developed with a new release every year. PostgreSQL supports the most advanced features included in SQL standards. It also provides NoSQL capabilities and very rich data types and extensions. All of this makes PostgreSQL a very attractive solution in software systems. If you run a database, you want it to perform well and you want to be able to secure it. As the world's most advanced open source database, PostgreSQL has unique built-in ways to achieve these goals. This book will show you a multitude of ways to enhance your database's performance and give you insights into measuring and optimizing a PostgreSQL database to achieve better performance. This book is your one-stop guide to elevate your PostgreSQL knowledge to the next level. First, you'll get familiarized with essential developer/administrator concepts such as load balancing, connection pooling, and distributing connections to multiple nodes. Next, you will explore memory optimization techniques before exploring the security controls offered by PostgreSQL. Then, you will move on to the essential database/server monitoring and replication strategies with PostgreSQL. Finally, you will learn about query processing algorithms. Style and approach This comprehensive guide is packed with practical administration tasks. Each topic is explained using examples and a step-by-step approach. A comprehensive series of dependable recipes to design, build, and implement a PostgreSQL server architecture free of common pitfalls that can operate for years to come. Each chapter is packed with instructions and examples to simplify even highly complex database operations. If you are a PostgreSQL DBA working on Linux systems who want a database that never gives up, this book is for you. If you've ever experienced a database outage, restored from a backup, spent hours trying to repair a malfunctioning cluster, or simply want to guarantee system stability, this book is definitely for you. Get to grips with building reliable, scalable, and maintainable database solutions for enterprises and production databases Key Features: Implement PostgreSQL 13 features to perform end-to-end modern database management Design, manage, and build enterprise database solutions using a unique recipe-based approach Solve common and not-so-common challenges faced while working to achieve optimal database performance Book Description: PostgreSQL has become the most advanced open source database on the market. This book follows a step-by-step approach, guiding you effectively in deploying PostgreSQL in production environments. The book starts with an introduction to PostgreSQL and its architecture. You'll cover common and not-so-common challenges faced while designing and managing the database. Next, the book focuses on backup and recovery strategies to ensure your database is steady and achieves optimal performance. Throughout the book, you'll address key challenges such as maintaining reliability, data integrity, a

fault-tolerant environment, a robust feature set, extensibility, consistency, and authentication. Moving ahead, you'll learn how to manage a PostgreSQL cluster and explore replication features for high availability. Later chapters will assist you in building a secure PostgreSQL server, along with covering recipes for encrypting data in motion and data at rest. Finally, you'll not only discover how to tune your database for optimal performance but also understand ways to monitor and manage maintenance activities, before learning how to perform PostgreSQL upgrades during downtime. By the end of this book, you'll be well-versed with the essential PostgreSQL 13 features to build enterprise relational databases.

**What You Will Learn:** Understand logical and physical backups in Postgres Demonstrate the different types of replication methods possible with PostgreSQL today Set up a high availability cluster that provides seamless automatic failover for applications Secure a PostgreSQL encryption through authentication, authorization, and auditing Analyze the live and historic activity of a PostgreSQL server Understand how to monitor critical services in Postgres 13 Manage maintenance activities and performance tuning of a PostgreSQL cluster Who this book is for: This PostgreSQL book is for database architects, database developers and administrators, or anyone who wants to become well-versed with PostgreSQL 13 features to plan, manage, and design efficient database solutions. Prior experience with the PostgreSQL database and SQL language is expected.

PostgreSQL offers a comprehensive set of replication related features. Unleashing the power of PostgreSQL provides you with countless opportunities and a competitive advantage over other database systems. This book will guide you through the most important concepts of PostgreSQL replication. It contains all the information you need to design and operate replicated setups. Beginning by giving you an understanding of replication concepts, the PostgreSQL transaction log, and Point-in-time Recovery, we gradually move on to setting up asynchronous and synchronous replication. Next up, you will learn to monitor a PostgreSQL cluster setup, deal with monitoring tools, and then move on to understanding Linux High Availability. Further, we explore widely-used tools such as Slony, SkyTools, Postgres-XC, and walbouncer, and set up PL/Proxy. Finally, you'll get acquainted with the new technology of BDR, which allows bidirectional replication in PostgreSQL.

This book will get you up and running with building efficient relational database solutions right from scratch with the newest features of PostgreSQL 11. You will learn the end-to-end working of relational databases and how to work with database structures. You will also be able to write essential SQL statements, perform data manipulation and ... Updated to include the new features introduced in PostgreSQL 13, this book shows you how to build better PostgreSQL applications and administer your PostgreSQL database efficiently. You'll master the advanced features of PostgreSQL and develop the skills you need to build secure and highly available database solutions.

PostgreSQL is increasingly utilized in all kind of applications, starting from desktop to web and mobile applications. In this book, you will find the best ways to design, monitor and maintain your PostgreSQL solution, with suggestions and tips for high performance, troubleshooting and high availability.

Master over 100 recipes to design and implement a highly available server with the advanced features of PostgreSQL  
About This Book Create a PostgreSQL cluster that stays online even when disaster strikes Avoid costly downtime and data loss that can ruin your business Updated to include the newest features introduced in PostgreSQL 9.6 with hands-on industry-driven recipes Who This Book Is For If you are a PostgreSQL DBA working on Linux systems who want a database that never gives up, this book is for you. If you've ever experienced a database outage, restored from a backup, spent hours trying to repair a malfunctioning cluster, or simply want to guarantee system stability, this book is definitely for you. What You Will Learn Protect your data with PostgreSQL replication and management tools such as Slony, Bucardo, pglogical, and WAL-E Hardware planning to help your database run efficiently Prepare for catastrophes and prevent them before they happen Reduce database resource contention with connection pooling using pgpool and PgBouncer Automate monitoring and alerts to visualize cluster activity using Nagios and collected Construct a robust software stack that can detect and fix outages Learn simple PostgreSQL High Availability with Patroni, or dive into the full power of Pacemaker. In Detail Databases are nothing without the data they store. In the event of a failure - catastrophic or otherwise - immediate recovery is essential. By carefully combining multiple servers, it's even possible to hide the fact a failure occurred at all. From hardware selection to software stacks and horizontal scalability, this book will help you build a versatile PostgreSQL cluster that will survive crashes, resist data corruption, and grow smoothly with customer demand. It all begins with hardware selection for the skeleton of an efficient PostgreSQL database cluster. Then it's on to preventing downtime as well as troubleshooting some real life problems that administrators commonly face. Next, we add database monitoring to the stack, using collectd, Nagios, and Graphite. And no stack is complete without replication using multiple internal and external tools, including the newly released pglogical extension. Pacemaker or Raft consensus tools are the final piece to grant the cluster the ability to heal itself. We even round off by tackling the complex problem of data scalability. This book exploits many new features introduced in PostgreSQL 9.6 to make the database more efficient and adaptive, and most importantly, keep it running. Style and approach This book contains practical recipes that will help the reader solve real world problems related to high availability in PostgreSQL. Every recipe is explained in detail, with relevant explanations, tips and tricks provided for quicker and easier understanding.

Master the capabilities of PostgreSQL 9.6 to efficiently manage and maintain your database About This Book Your one-stop guide to mastering the advanced concepts in PostgreSQL with ease Master query optimization, replication, and high

availability with PostgreSQL Extend the functionalities of PostgreSQL to suit your organizational needs with minimum effort Who This Book Is For If you are a PostgreSQL data architect or an administrator who wants to understand how to implement advanced functionalities and master complex administrative tasks with PostgreSQL, then this book is perfect for you. Prior experience of administrating a PostgreSQL database and a working knowledge of SQL is required to make the best use of this book. What You Will Learn Get to grips with the advanced features of PostgreSQL 9.6 and handle advanced SQL Make use of the indexing features in PostgreSQL and fine-tune the performance of your queries Work with the stored procedures and manage backup and recovery Master the replication and failover techniques Troubleshoot your PostgreSQL instance for solutions to the common and not-so-common problems Learn how to migrate your database from MySQL and Oracle to PostgreSQL without any hassle In Detail PostgreSQL is an open source database used for handling large datasets (Big Data) and as a JSON document database. It also has applications in the software and web domains. This book will enable you to build better PostgreSQL applications and administer databases more efficiently. We begin by explaining the advanced database design concepts in PostgreSQL 9.6, along with indexing and query optimization. You will also see how to work with event triggers and perform concurrent transactions and table partitioning, along with exploring SQL and server tuning. We will walk you through implementing advanced administrative tasks such as server maintenance and monitoring, replication, recovery and high availability, and much more. You will understand the common and not-so-common troubleshooting problems and how you can overcome them. By the end of this book, you will have an expert-level command of the advanced database functionalities and will be able to implement advanced administrative tasks with PostgreSQL. Style and Approach This book is a comprehensive guide covering all the concepts you need to master PostgreSQL. Packed with hands-on examples, tips and tricks, even the most advanced concepts are explained in a very easy-to-follow manner. Every chapter in the book does not only focus on how each task is performed, but also why.

Obtain all the skills you need to configure and manage a PostgreSQL database. In this book you will begin by installing and configuring PostgreSQL on a server by focusing on system-level parameter settings before installation. You will also look at key post-installation steps to avoid issues in the future. The basic configuration of PostgreSQL is tuned for compatibility rather than performance. Keeping this in mind, you will fine-tune your PostgreSQL parameters based on your environment and application behavior. You will then get tips to improve database monitoring and maintenance followed by database security for handling sensitive data in PostgreSQL. Every system containing valuable data needs to be backed-up regularly. PostgreSQL follows a simple back-up procedure and provides fundamental approaches to back up your data. You will go through these approaches and choose the right one based on your environment. Running your

application with limited resources can be tricky. To achieve this you will implement a pooling mechanism for your PostgreSQL instances to connect to other databases. Finally, you will take a look at some basic errors faced while working with PostgreSQL and learn to resolve them in the quickest manner. What You Will Learn Configure PostgreSQL for performance Monitor and maintain PostgreSQL instances Implement a backup strategy for your data Resolve errors faced while using PostgreSQL Who This Book Is For Readers with basic knowledge of PostgreSQL who wish to implement key solutions based on their environment.

The one-stop-source powering High Availability success, jam-packed with ready to use insights for results, loaded with all the data you need to decide how to gain and move ahead. Based on extensive research, this lays out the thinking of the most successful High Availability knowledge experts, those who are adept at continually innovating and seeing opportunities. This is the first place to go for High Availability innovation - INCLUDED are numerous real-world High Availability blueprints, presentations and templates ready for you to access and use. Also, if you are looking for answers to one or more of these questions then THIS is the title for you: High Availability: How clustering multiple load balancers ? High Availability: What is a canary request? What is High Availability and why is it important? How can I learn more about programming for high-availability clusters with PHP? NoSQL: What database should I use to get the best high availability/redundancy/uptime? What are some good open source stacks for building a high-availability Complex Event Processing system? How do I test high-availability configuration? High Availability: How to setup NLB in Oracle 11g RAC? How can SDN be applied in the context of disaster recovery and high availability? Does CDH 5 contain any private security or High Availability (HA) features? How do I achieve high availability in a cloud using Openstack? How do I configure a PostgreSQL high-availability cluster? What are the best sites to discuss about high availability architectures for different use cases? What Do You Mean By High Availability? What is the easiest way to ensure high availability using active replication? What are users' experiences with C# Mono in large scale, high availability environments? What's the best methodology for providing high availability for a Subversion repository? What are case scenarios for BMC Remedy products with Oracle High Availability ? ...and much more..."

Over 100 recipes to design and implement a highly available server with the advanced features of PostgreSQL 9.4,9.5 and 9.6About This Book\* Create a PostgreSQL cluster that stays online even when disaster strikes\* Avoid costly downtime and data loss that can ruin your business\* Updated to include the newest features introduced in PostgreSQL 9.6 with hands-on industry-driven recipesWho This Book Is ForIf you are a PostgreSQL DBA working on Linux systems who want a database that never gives up, this book is for you. If you've ever experienced a database outage, restored from a backup, spent hours trying to repair a malfunctioning cluster, or simply want to guarantee system stability, this book is definitely for you.What you will learn\* Protect your

data with PostgreSQL replication and management tools such as Slony, Bucardo, pglogical, and WAL-E\* Hardware planning to help your database run efficiently\* Prepare for catastrophes and prevent them before they happen\* Reduce database resource contention with connection pooling using pgpool and PgBouncer\* Automate monitoring and alerts to visualize cluster activity using Nagios and collected\* Construct a robust software stack that can detect and fix outages\* Learn simple PostgreSQL High Availability with Patroni, or dive into the full power of Pacemaker. In Detail Databases are nothing without the data they store. In the event of a failure - catastrophic or otherwise - immediate recovery is essential. By carefully combining multiple servers, it's even possible to hide the fact a failure occurred at all. From hardware selection to software stacks and horizontal scalability, this book will help you build a versatile PostgreSQL cluster that will survive crashes, resist data corruption, and grow smoothly with customer demand. It all begins with hardware selection for the skeleton of an efficient PostgreSQL database cluster. Then it's on to preventing downtime as well as troubleshooting some real life problems that administrators commonly face. Next, we add database monitoring to the stack, using collectd, Nagios, and Graphite. And no stack is complete without replication using multiple internal and external tools, including the newly released pglogical extension. Pacemaker or Raft consensus tools are the final piece to grant the cluster the ability to heal itself. We even round off by tackling the complex problem of data scalability. This book exploits many new features introduced in PostgreSQL 9.6 to make the database more efficient and adaptive, and most importantly, keep it running.

Filled with practical, step-by-step instructions and clear explanations for the most important and useful tasks. This hands-on guide provides a quick and easy way to back up and restore your database using PostgreSQL. Written for database administrators who want to create backups of their critical enterprise data and efficiently restore it using PostgreSQL.

A practical guide to administer, monitor and replicate your PostgreSQL 11 database Key Features Study and apply the newly introduced features in PostgreSQL 11 Tackle any problem in PostgreSQL 11 administration and management Catch up on expert techniques for monitoring, fine-tuning, and securing your database Book Description PostgreSQL is a powerful, open source database management system with an enviable reputation for high performance and stability. With many new features in its arsenal, PostgreSQL 11 allows you to scale up your PostgreSQL infrastructure. This book takes a step-by-step, recipe-based approach to effective PostgreSQL administration. The book will introduce you to new features such as logical replication, native table partitioning, additional query parallelism, and much more to help you to understand and control, crash recovery and plan backups. You will learn how to tackle a variety of problems and pain points for any database administrator such as creating tables, managing views, improving performance, and securing your database. As you make steady progress, the book will draw attention to important topics such as monitoring roles, backup, and recovery of your PostgreSQL 11 database to help you understand roles and produce a summary of log files, ensuring high availability, concurrency, and replication. By the end of this book, you will have the necessary knowledge to manage your PostgreSQL 11 database efficiently. What you will learn Troubleshoot open source PostgreSQL version 11 on various platforms Deploy best practices for planning and designing live databases Select and

implement robust backup and recovery techniques in PostgreSQL 11 Use pgAdmin or OmniDB to perform database administrator (DBA) tasks Adopt efficient replication and high availability techniques in PostgreSQL Improve the performance of your PostgreSQL solution Who this book is for This book is designed for database administrators, data architects, database developers, or anyone with an interest in planning and running live production databases using PostgreSQL 11. It is also ideal if you're looking for hands-on solutions to any problem associated with PostgreSQL 11 administration. Some experience with handling PostgreSQL databases will be beneficial

Learning Heroku Postgres is targeted at developers and database admins. Even if you're new to Heroku Postgres, you'll be able to master both the basic as well as advanced features of Heroku Postgres. Since Heroku Postgres is incredibly user-friendly, no previous experience in computer coding or programming is required.

Various High Availability DataBase systems (HADB) are used to provide high availability. Pairing an active database system with a standby system is one commonly used HADB techniques. The active system serves read/write workloads. One or more standby systems replicate the active and serve read-only workloads. Though widely used, this technique has some significant drawbacks: The active system becomes the bottleneck under heavy write workloads. Replicating changes synchronously from the active to the standbys further reduces the performance of the active system. Asynchronous replication, however, risk the loss of updates during failover. The shared-nothing architecture of active-standby systems is unnecessarily complex and cost inefficient. In this thesis we present SHADOW systems, a new technique for database high availability. In a SHADOW system, the responsibility for database replication is pushed from the database systems into a shared, reliable, storage system. The active and standby systems share access to a single logical copy of the database, which resides in shared storage. SHADOW introduces write offloading, which frees the active system from the need to update the persistent database, placing that responsibility on the underutilized standby system instead. By exploiting shared storage, SHADOW systems avoid the overhead of database-managed synchronized replication, while ensuring that no updates will be lost during a failover. We have implemented a SHADOW system using PostgreSQL, and we present the results of a performance evaluation that shows that the SHADOW system can outperform both traditional synchronous replication and standalone PostgreSQL systems.

Design, develop and deploy a highly available vSphere environment for VMware Horizon View About This Book Enhance your capability of meeting various Service Level Agreements in VMware Horizon View Get acquainted through all the necessary considerations for building a View environment Cover VMware High Availability hurdle by hurdle along with the checklists for verification of the environment being ready for production Who This Book Is For If you manage, plan or deploy VMware Horizon View or are looking for tips for best practices and configuration details this book is for you. This book is intended for administrators who design and deploy VMware Horizon View or administrators who are looking for ways to improve their existing View environment. What You Will Learn Install and configure a VMware Horizon View Connection Server and redundant pair Discover the networking requirements for View and learn how to build redundancy into your network Analyze each of the View user pool

types and how each one can be made highly available and survivable. Get to know about storage protocols such as NFS, iSCSI and Fibre Channel Deploy Virtual SAN, and find out how to effectively couple Virtual SAN with View Learn about View monitoring tools to allow fast responses to various crises Plan, analyze and upgrade VMware Horizon View Analyze network services required for VMware Horizon View and build them in a redundant manner In Detail The increasing movement to virtualize workloads and workstations has put VMware Horizon View into a central mission critical role in many environments. Administrators may be overwhelmed with planning for outages and dealing with failure scenarios. It's easy to miss small details that will result in outages down the road. Following VMware Horizon View best practices and planning ahead with network infrastructure will allow you to avoid these common pit falls. This book will walk you through the setup and configuration of View in a highly available configuration. It will provide you with the skills to analyze and deploy configurations that can stand up to rigorous failure standards. The book starts with deploying and basic configuration of VMware Horizon View in a redundant setup, then moves on to cover high availability for networking, fibre channel, NFS, and iSCSI. We finish this book with monitoring and upgrade planning. At the end we also learn about maintaining the uptime and minimizing the downtime that can be caused due to various factors. Each topic comes with a list of best practices and failure scenarios to test. Administrators will learn the intricacies of protecting a View environment. Style and approach This book provides configuration and installation steps for administration and installation of a Horizon View server. It includes high-level overviews of any protocols, services used by Horizon View, and best practices and high availability checklists for each chapter.

A comprehensive guide to understanding key techniques for architecture and hardware planning, monitoring, replication, backups, and decoupling Key Features Newly updated edition, covering the latest PostgreSQL 12 features with hands-on industry-driven recipes Create a PostgreSQL cluster that stays online even when disaster strikes Learn how to avoid costly downtime and data loss that can ruin your business Book Description Databases are nothing without the data they store. In the event of an outage or technical catastrophe, immediate recovery is essential. This updated edition ensures that you will learn the important concepts related to node architecture design, as well as techniques such as using repmgr for failover automation. From cluster layout and hardware selection to software stacks and horizontal scalability, this PostgreSQL cookbook will help you build a PostgreSQL cluster that will survive crashes, resist data corruption, and grow smoothly with customer demand. You'll start by understanding how to plan a PostgreSQL database architecture that is resistant to outages and scalable, as it is the scaffolding on which everything rests. With the bedrock established, you'll cover the topics that PostgreSQL database administrators need to know to manage a highly available cluster. This includes configuration, troubleshooting, monitoring and alerting, backups through proxies, failover automation, and other considerations that are essential for a healthy PostgreSQL cluster. Later, you'll learn to use multi-master replication to maximize server availability. Later chapters will guide you through managing major version upgrades without downtime. By the end of this book, you'll have learned how to build an efficient and adaptive PostgreSQL 12 database cluster. What you will learn Understand how to protect data with PostgreSQL replication tools Focus on hardware planning to ensure that

your database runs efficiently Reduce database resource contention with connection pooling Monitor and visualize cluster activity with Nagios and the TIG (Telegraf, InfluxDB, Grafana) stack Construct a robust software stack that can detect and avert outages Use multi-master to achieve an enduring PostgreSQL cluster Who this book is for This book is for Postgres administrators and developers who are looking to build and maintain a highly reliable PostgreSQL cluster. Although knowledge of the new features of PostgreSQL 12 is not required, a basic understanding of PostgreSQL administration is expected.

A practical guide to administer, monitor and replicate your PostgreSQL 10 database Key Features Get to grips with the capabilities of PostgreSQL 10 to administer your database more efficiently Monitor, tune, secure and protect your database for optimal performance A step-by-step, recipe-based guide to help you tackle any problem in PostgreSQL 10 administration with ease Book Description PostgreSQL is a powerful, open source database management system with an enviable reputation for high performance and stability. With many new features in its arsenal, PostgreSQL 10 allows users to scale up their PostgreSQL infrastructure. This book takes a step-by-step, recipe-based approach to effective PostgreSQL administration. Throughout this book, you will be introduced to these new features such as logical replication, native table partitioning, additional query parallelism, and much more. You will learn how to tackle a variety of problems that are basically the pain points for any database administrator - from creating tables to managing views, from improving performance to securing your database. More importantly, the book pays special attention to topics such as monitoring roles, backup, and recovery of your PostgreSQL 10 database, ensuring high availability, concurrency, and replication. By the end of this book, you will know everything you need to know to be the go-to PostgreSQL expert in your organization. What you will learn Get to grips with the newly released PostgreSQL 10 features to improve database performance and reliability Manage open source PostgreSQL versions 10 on various platforms. Explore best practices for planning and designing live databases Select and implement robust backup and recovery techniques in PostgreSQL 10 Explore concise and clear guidance on replication and high availability Discover advanced technical tips for experienced users Who this book is for This book is for database administrators, data architects, developers, or anyone with an interest in planning for, or running, live production databases using PostgreSQL. It is most suited to those looking for hands-on solutions to any problem associated with PostgreSQL administration.

If you are a system administrator, database administrator, architect, developer, or anyone with an interest in planning, managing, and designing database solutions using PostgreSQL, this is the book for you. This book is suited for you if you have some prior experience with any relational database or with the SQL language.

A guide to building applications with Rails covers such topics as metaprogramming, Active Support library, advanced database functions, security principles, RESTful architecture, and optimizing performance.

Get started with PostgreSQL on the cloud and discover the advantages, disadvantages, and limitations of the cloud services from Amazon, Rackspace, Google, and Azure. Once you have chosen your cloud service, you will focus on

securing it and developing a back-up strategy for your PostgreSQL instance as part of your long-term plan. Beginning PostgreSQL on the Cloud covers other essential topics such as setting up replication and high availability; encrypting your saved cloud data; creating a connection pooler for your database; and monitoring PostgreSQL on the cloud. The book concludes by showing you how to install and configure some of the tools that will help you get started with PostgreSQL on the cloud. This book shows you how database as a service enables you to spread your data across multiple data centers, ensuring that it is always accessible. You'll discover that this model does not expect you to install and maintain databases yourself because the database service provider does it for you. You no longer have to worry about the scalability and high availability of your database. What You Will Learn Migrate PostgreSQL to the cloud Choose the best configuration and specifications of cloud instances Set up a backup strategy that enables point-in-time recovery Use connection pooling and load balancing on cloud environments Monitor database environments on the cloud Who This Book Is For Those who are looking to migrate to PostgreSQL on the Cloud. It will also help database administrators in setting up a cloud environment in an optimized way and help them with their day-to-day tasks. Discover the methodologies and best practices for getting started with Google Cloud Platform relational services – CloudSQL and CloudSpanner. The book begins with the basics of working with the Google Cloud Platform along with an introduction to the database technologies available for developers from Google Cloud. You'll then take an in-depth hands on journey into Google CloudSQL and CloudSpanner, including choosing the right platform for your application needs, planning, provisioning, designing and developing your application. Sample applications are given that use Python to connect to CloudSQL and CloudSpanner, along with helpful features provided by the engines. You'll also implement practical best practices in the last chapter. Hands On Google Cloud SQL and Cloud Spanner is a great starting point to apply GCP data offerings in your technology stack and the code used allows you to try out the examples and extend them in interesting ways. What You'll Learn Get started with Big Data technologies on the Google Cloud Platform Review CloudSQL and Cloud Spanner from basics to administration Apply best practices and use Google's CloudSQL and CloudSpanner offering Work with code in Python notebooks and scripts Who This Book Is For Application architects, database architects, software developers, data engineers, cloud architects. Cloud services are just as susceptible to network outages as any other platform. This concise book shows you how to prepare for potentially devastating interruptions by building your own resilient and reliable applications in the public cloud. Guided by engineers from 9apps—an independent provider of Amazon Web Services and Eucalyptus cloud solutions—you'll learn how to combine AWS with open source tools such as PostgreSQL, MongoDB, and Redis. This isn't a book on theory. With detailed examples, sample scripts, and solid advice, software engineers with operations

experience will learn specific techniques that 9apps routinely uses in its cloud infrastructures. Build cloud applications with the "rip, mix, and burn" approach Get a crash course on Amazon Web Services Learn the top ten tips for surviving outages in the cloud Use elasticsearch to build a dependable NoSQL data store Combine AWS and PostgreSQL to build an RDBMS that scales well Create a highly available document database with MongoDB Replica Set and SimpleDB Augment Redis with AWS to provide backup/restore, failover, and monitoring capabilities Work with CloudFront and Route 53 to safeguard global content delivery

Get up to speed with core PostgreSQL tasks such as database administration, application development, database performance monitoring, and database testing Key Features Build real-world enterprise database management systems using Postgres 12 features Explore the development, administrative and security aspects of PostgreSQL 12 Implement best practices from industry experts to build powerful database applications Book Description PostgreSQL is an open-source object-relational database management system (DBMS) that provides enterprise-level services, including high performance and scalability. This book is a collection of unique projects providing you with a wealth of information relating to administering, monitoring, and testing PostgreSQL. The focus of each project is on both the development and the administrative aspects of PostgreSQL. Starting by exploring development aspects such as database design and its implementation, you'll then cover PostgreSQL administration by understanding PostgreSQL architecture, PostgreSQL performance, and high-availability clusters. Various PostgreSQL projects are explained through current technologies such as DevOps and cloud platforms using programming languages like Python and Node.js. Later, you'll get to grips with the well-known database API tool, PostgREST, before learning how to use popular PostgreSQL database testing frameworks. The book is also packed with essential tips and tricks and common patterns for working seamlessly in a production environment. All the chapters will be explained with the help of a real-world case study on a small banking application for managing ATM locations in a city. By the end of this DBMS book, you'll be proficient in building reliable database solutions as per your organization's needs. What you will learn Set up high availability PostgreSQL database clusters in the same containment, a cross-containment, and on the cloud Monitor the performance of a PostgreSQL database Create automated unit tests and implement test-driven development for a PostgreSQL database Develop PostgreSQL apps on cloud platforms using DevOps with Python and Node.js Write robust APIs for PostgreSQL databases using Python programming, Node.js, and PostgREST Create a geospatial database using PostGIS and PostgreSQL Implement automatic configuration by Ansible and Terraform for Postgres Who this book is for This PostgreSQL book is for database developers, database administrators, data architects, or anyone who wants to build end-to-end database projects using Postgres. This book will also appeal to software engineers, IT technicians, computer

science researchers, and university students who are interested in database development and administration. Some familiarity with PostgreSQL and Linux is required to grasp the concepts covered in the book effectively.

A comprehensive guide to understanding key techniques for architecture and hardware planning, monitoring, replication, backups, and decoupling Key Features Newly updated edition, covering the latest PostgreSQL 12 features with hands-on industry-driven recipes Create a PostgreSQL cluster that stays online even when disaster strikes Learn how to avoid costly downtime and data loss that can ruin your business Book Description Databases are nothing without the data they store. In the event of an outage or technical catastrophe, immediate recovery is essential. This updated edition ensures that you will learn the important concepts related to node architecture design, as well as techniques such as using repmgr for failover automation. From cluster layout and hardware selection to software stacks and horizontal scalability, this PostgreSQL cookbook will help you build a PostgreSQL cluster that will survive crashes, resist data corruption, and grow smoothly with customer demand. You'll start by understanding how to plan a PostgreSQL database architecture that is resistant to outages and scalable, as it is the scaffolding on which everything rests. With the bedrock established, you'll cover the topics that PostgreSQL database administrators need to know to manage a highly available cluster. This includes configuration, troubleshooting, monitoring and alerting, backups through proxies, failover automation, and other considerations that are essential for a healthy PostgreSQL cluster. Later, you'll learn to use multi-master replication to maximize server availability. Later chapters will guide you through managing major version upgrades without downtime. By the end of this book, you'll have learned how to build an efficient and adaptive PostgreSQL 12 database cluster. What you will learn Understand how to protect data with PostgreSQL replication tools Focus on hardware planning to ensure that your database runs efficiently Reduce database resource contention with connection pooling Monitor and visualize cluster activity with Nagios and the TIG (Telegraf, InfluxDB, Grafana) stack Construct a robust software stack that can detect and avert outages Use multi-master to achieve an enduring PostgreSQL cluster Who this book is for This book is for Postgres administrators and developers who are looking to build and maintain a highly reliable PostgreSQL cluster. Although knowledge of the new features of PostgreSQL 12 is not required, a basic understanding of PostgreSQL administration is expected.

[Copyright: 48788b4021bcf17724692d84a2260294](#)