

Chitij Chauhan, Dinesh Kumar

PostgreSQL High Performance Cookbook

Mastering query optimization, database monitoring, and performance-tuning for PostgreSQL



Packt>

PostgreSQL High Performance Cookbook

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. Also It also provides NoSQL capabilities and very rich data types and extensions. All of this makes PostgreSQL a very attractive solution in various kinds of software systems.

If you run a database, you want good performance, and you want to be able to secure it as much as possible. PostgreSQL, as the world's most advanced open source database, has built-in unique ways to achieve these goals. This book will show you plenty of ways to enhance your database's performance and gives you insights into measuring and optimizing a PostgreSQL database to achieve better performance.

This book is your one-stop guide taking your PostgreSQL knowledge to the next level. We'll walk through essential developer/administrator concepts such as load balancing, connection pooling, and distributing connections to multiple nodes. Next, we will explore various memory optimization techniques before exploring the security controls offered by PostgreSQL. Later, you will move on to essential database/server monitoring and replication strategies with PostgreSQL. Finally, you will learn about various query processing algorithms.

Things you will learn:

- Build replication strategies for homogeneous and heterogeneous databases
- Test and build a secure PostgreSQL database with multiple benchmarking techniques
- Find out how to manage the replication using multiple tools
- Find out how to monitor the database and operating system metrics
- 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

Packt
www.packtpub.com

\$ 54.99 US
£ 45.99 UK

Prices do not include local sales
Tax or VAT where applicable



PostgreSQL High Performance Cookbook

Chitij Chaudhan, Dinesh Kumar



PostgreSQL High Performance Cookbook

Mastering query optimization, database monitoring, and performance-tuning for PostgreSQL

Chitij Chauhan
Dinesh Kumar



BIRMINGHAM - MUMBAI

PostgreSQL High Performance Cookbook

Copyright © 2017 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the authors, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: March 2017

Production reference: 1240317

Published by Packt Publishing Ltd.

Livery Place

35 Livery Street

Birmingham

B3 2PB, UK.

ISBN 978-1-78528-433-5

www.packtpub.com

Credits

Authors

Chitij Chauhan
Dinesh Kumar

Copy Editor

Safis Editing

Reviewers

Baji Shaik
Feng Tan

Project Coordinator

Nidhi Joshi

Commissioning Editor

Dipika Gaonkar

Proofreader

Safis Editing

Acquisition Editor

Nitin Dasan

Indexer

Aishwarya Gangawane

Content Development Editor

Mayur Pawanikar

Production Coordinator

Shraddha Falebhai

Technical Editor

Dinesh Pawar

About the Authors

Chitij Chauhan currently works as a senior database administrator at an IT-based MNC in Chandigarh. He has over 10 years of work experience in the field of database and system administration, with specialization in MySQL clustering, PostgreSQL, Greenplum, Informix DB2, SQL Server 2008, Sybase, and Oracle. He is a leading expert in the area of database security, with expertise in database security products such as IBM InfoSphere Guardium, Oracle Database Vault, and Imperva.

Dinesh Kumar is an enthusiastic open source developer and has written several open source tools for PostgreSQL. He recently announced pgBucket, a brand new job scheduler for PostgreSQL. He is also a frequent blogger at manojadinesh.blogspot.com, where he talks more about PostgreSQL. He is currently working as a senior database engineer in OpenSCG and building the PostgreSQL cloud operations. He has more than 6 years of experience as an Oracle and PostgreSQL database administrator and developer, and is currently focusing on PostgreSQL.

Thanks to my loving parents, Sreenivasulu and Vanamma, who raised me in a small village called Viruvuru, which is in the Nellore district of India. Thanks to my loving wife, Manoja, who enlightens my life with her wonderful support. Also, thanks to my friend Baji Shaik and coordinators, Mayur Pawanikar and Nitin Dasan, for their excellent support. Finally, thanks to every PostgreSQL contributor, author, and blogger.

About the Reviewers

Baji Shaik is a database administrator and developer. He is a co-author of *PostgreSQL Development Essentials* and has tech-reviewed *Troubleshooting PostgreSQL* by Packt Publishing. He is currently working as a database consultant at OpenSCG. He has an engineering degree in telecommunications, and had started his career as a C# and Java developer. Baji started working with databases in 2011, and over the years he has worked with Oracle, PostgreSQL, and Greenplum. His background spans a length and breadth of expertise and experience in SQL/NoSQL database technologies. He has good knowledge of automation, orchestration, and DevOps in a cloud environment. He likes to watch movies, read books, and write technical blogs. He also loves to spend time with family. Baji is a certified PostgreSQL professional.

Feng Tan is from China. His nickname is Francs. He was a PostgreSQL DBA at SkyMobi (NASDAQ: MOBI) for more than 5 years, where he was maintaining more than 100 PostgreSQL instances. He gave presentations at the China PostgreSQL conference on topics such as Oracle VS PostgreSQL and PostgreSQL 9.4 new features.

Feng Tan likes to share PostgreSQL technology in his blog at <http://francs3.blog.163.com/>. He is also one of the translators of PostgreSQL 9 Administration Cookbook Chinese Edition.

Currently, he serves as the open source database administrator at China Mobile Group Zhejiang Co. Ltd.

www.PacktPub.com

For support files and downloads related to your book, please visit www.PacktPub.com.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.PacktPub.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at service@packtpub.com for more details.

At www.PacktPub.com, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.



<https://www.packtpub.com/mapt>

Get the most in-demand software skills with Mapt. Mapt gives you full access to all Packt books and video courses, as well as industry-leading tools to help you plan your personal development and advance your career.

Why subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print, and bookmark content
- On demand and accessible via a web browser

Customer Feedback

Thanks for purchasing this Packt book. At Packt, quality is at the heart of our editorial process. To help us improve, please leave us an honest review on this book's Amazon page at <https://www.amazon.com/dp/1785284339>.

If you'd like to join our team of regular reviewers, you can e-mail us at customerreviews@packtpub.com. We award our regular reviewers with free eBooks and videos in exchange for their valuable feedback. Help us be relentless in improving our products!

Table of Contents

Preface	1
Chapter 1: Database Benchmarking	7
Introduction	7
CPU benchmarking	8
Getting ready	8
How to do it...	8
Phoronix	8
sysbench	10
How it works...	10
Phoronix	10
sysbench	10
Memory benchmarking	11
Getting ready	11
How to do it...	11
Phoronix	11
tmpfs	11
Write test	12
Read test	13
How it works...	13
Disk benchmarking	13
Getting ready	13
How to do it...	13
Phoronix	14
bonnie++	14
How it works...	15
bonnie++	15
Performing a seek rate test	15
Getting ready	16
How to do it...	16
How it works...	16
Working with the fsync commit rate	17
Getting ready	17
How to do it...	17
How it works...	17
Checking IOPS	19
Getting ready	19

How to do it...	19
Sequential mixed read and write	19
Random mixed read and write	20
How it works...	20
Storage sizing	21
Getting ready	21
How to do it...	22
How it works...	22
Discussing RAID levels	22
Getting ready	22
How to do it...	23
How it works...	23
RAID 0	23
RAID 1	24
RAID 5	24
RAID 6	24
RAID 10	24
Configuring pgbench	25
Getting ready	25
How to do it...	25
How it works...	26
Running read/write pgbench test cases	27
Getting ready	27
How to do it...	27
Read-only	27
Write-only	27
How it works...	28
Chapter 2: Server Configuration and Control	29
Introduction	29
Starting the server manually	29
Getting ready	30
How to do it...	30
How it works...	30
Stopping the server quickly	31
Getting ready	31
How to do it...	32
How it works...	32
Stopping the server in an emergency	32
How to do it...	32
How it works...	32

Reloading server configuration	33
Getting ready	33
How to do it...	33
How it works...	33
Restarting the database server quickly	34
How to do it...	34
How it works...	34
Tuning connection-related parameters	35
How to do it...	35
How it works...	35
Tuning query-related parameters	36
How to do it...	37
How it works...	37
Tuning logging-related parameters	39
How to do it...	39
How it works...	39
Chapter 3: Device Optimization	41
Introduction	41
Understanding memory units in PostgreSQL	41
Getting ready	41
How to do it...	42
shared_buffers	42
temp_buffers	42
work_mem	42
maintenance_work_mem	42
wal_buffers	43
max_stack_depth	43
effective_cache_size	43
How it works...	43
Handling Linux/Unix memory parameters	43
Getting ready	43
How to do it...	44
kernel.shmmax	44
kernel.shmall	44
kernel.shmmni	44
vm.swappiness	45
vm.overcommit_memory	45
vm.overcommit_ratio	45
vm.dirty_background_ratio	45
vm.dirty_ratio	46
How it works...	46

CPU scheduling parameters	46
Getting ready	46
How to do it...	47
kernel.sched_autogroup_enabled	47
kernel.sched_min_granularity_ns	47
kernel.sched_latency_ns	47
kernel.sched_wakeup_granularity_ns	47
kernel.sched_migration_cost_ns	48
How it works...	48
Disk tuning parameters	48
Getting ready	48
How to do it...	49
CFQ	49
noop	49
Deadline	49
How it works...	50
Identifying checkpoint overhead	50
Getting ready	50
How to do it...	50
How it works...	51
Analyzing buffer cache contents	51
Getting ready	51
How to do it...	52
How it works...	53
Chapter 4: Monitoring Server Performance	55
Introduction	55
Monitoring CPU usage	56
Getting ready	56
How to do it...	56
How it works...	56
Monitoring paging and swapping	57
Getting ready	57
How to do it...	58
How it works...	60
Tracking CPU consuming processes	60
Getting ready	60
How to do it...	60
How it works...	61
Monitoring CPU load	61
How to do it...	62

How it works...	62
Identifying CPU bottlenecks	62
How to do it...	62
How it works...	63
Identifying disk I/O bottlenecks	63
How to do it...	64
How it works...	65
Monitoring system load	65
How to do it...	65
How it works...	67
Tracking historical CPU usage	67
Getting ready	67
How to do it...	68
How it works...	68
There's more...	68
Tracking historical memory usage	69
Getting ready	69
How to do it...	69
How it works...	70
Monitoring disk space	70
How to do it...	70
How it works...	71
Monitoring network status	72
Getting ready	72
How to do it...	72
How it works...	72
Chapter 5: Connection Pooling and Database Partitioning	73
Introduction	73
Installing pgpool-II	74
Getting Ready	74
How to do it...	75
How it works...	76
Configuring pgpool and testing the setup	76
Getting ready	76
How to do it...	76
How it works...	77
There's more...	78
Installing PgBouncer	80
Getting ready	80

How to do it...	80
How it works...	82
Connection pooling using PgBouncer	82
Getting ready	82
How to do it...	82
How it works...	84
There's more...	85
Managing PgBouncer	86
Getting ready	86
How to do it...	86
How it works...	87
Implementing partitioning	89
Getting ready	89
How to do it...	90
How it works...	91
Managing partitions	93
Getting ready	94
How to do it...	94
How it works...	96
Installing PL/Proxy	97
How to do it...	98
How it works...	98
Partitioning with PL/Proxy	98
Getting ready	98
How to do it...	98
How it works...	102
Chapter 6: High Availability and Replication	105
Introduction	105
Setting up hot streaming replication	105
Getting ready	106
How to do it...	106
How it works...	109
Replication using Slony	110
Getting ready	111
How to do it...	111
How it works...	114
Replication using Londiste	115
Getting ready	115
How to do it...	116