

Shaun M. Thomas

PostgreSQL High Availability Cookbook

Second Edition

Master over 100 recipes to design and implement a highly available server with the advanced features of PostgreSQL



Packt>

PostgreSQL High Availability Cookbook

Second Edition

Master over 100 recipes to design and implement a highly available server with the advanced features of PostgreSQL

Shaun M. Thomas



BIRMINGHAM - MUMBAI

PostgreSQL High Availability Cookbook

Second Edition

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 author, 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: February 2017

Production reference: 1010217

Published by Packt Publishing Ltd.

Livery Place

35 Livery Street

Birmingham

B3 2PB, UK.

ISBN 978-1-78712-553-7

www.packtpub.com

Credits

Author

Shaun M. Thomas

Reviewer

Sheldon E. Strauch

Commissioning Editor

Amey Varangaonkar

Acquisition Editor

Ajith Menon

Content Development Editor

Sumeet Sawant

Technical Editor

Akash Patel

Copy Editors

Safis Editing

Project Coordinator

Shweta H Birwatkar

Proofreader

Safis Editing

Indexer

Rekha Nair

Graphics

Disha Haria

Production Coordinator

Melwyn Dsa

About the Author

Shaun M. Thomas has been working with PostgreSQL since late 2000. He is a frequent contributor to the PostgreSQL Performance and General mailing lists, assisting other DBAs with the knowledge he's gained over the years. In 2011 and 2012, he gave presentations at the Postgres Open conference on topics such as handling extreme throughput, high availability, server redundancy, and failover techniques. Most recently, he has contributed the Shard Manager extension and the walctl WAL management suite. Currently, he serves as the database architect at PEAK6 Investments, where he develops standard operating procedure (SOP) guidelines to facilitate reliable server architecture among many other tasks. Many of the techniques used in this book were developed specifically for this extreme environment. He believes that PostgreSQL has a stupendous future ahead, and he can't wait to see the advancements subsequent versions will bring.

<https://www.linkedin.com/in/bonesmoses>

About the Reviewer

Sheldon Strauch is a 22-year veteran of software consulting at companies such as IBM, Sears, Ernst & Young, and Kraft Foods. He has a Bachelor's degree in Business Administration and leverages his technical skills to improve the business' self-awareness. His interests include data gathering, management, and mining; maps and mapping; business intelligence; and application of data analysis for continuous improvement. He is currently focused on development of end-to-end data management and mining at Enova International, a financial services company located in Chicago. In his spare time, he enjoys the performing arts, particularly music, and traveling with his wife, Marilyn.

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

Thank you for purchasing this Packt book. We take our commitment to improving our content and products to meet your needs seriously - that's why your feedback is so valuable. Whatever your feelings about your purchase, please consider leaving a review on this book's Amazon page. Not only will this help us, more importantly it will also help others in the community to make an informed decision about the resources that they invest in to learn.

You can also review for us on a regular basis by joining our reviewers' club. **If you're interested in joining, or would like to learn more about the benefits we offer, please contact us:** customerreviews@packtpub.com.

Table of Contents

Preface	1
Chapter 1: Hardware Planning	9
Introduction	9
Planning for redundancy	10
Getting ready	10
How to do it...	11
How it works...	11
There's more...	12
See also	13
Having enough IOPS	13
Getting ready	13
How to do it...	14
How it works...	14
There's more...	15
A working example	16
Making concessions	16
Sizing storage	17
Getting ready	17
How to do it...	18
How it works...	19
There's more...	19
Real-world example	19
Adjusting the numbers	20
Incorporating the spreadsheet	21
Investing in a RAID	21
Getting ready	22
How to do it...	22
How it works...	23
There's more...	24
See also	24
Picking a processor	25
Getting ready	25
How to do it...	25
How it works...	27
There's more...	27

Hyperthreading	28
Turbo Boost	28
Power usage	28
See also	29
Making the most of memory	29
Getting ready	29
How to do it...	30
How it works...	30
There's more...	31
Exploring nimble networking	31
Getting ready	32
How to do it...	32
How it works...	34
There's more...	35
A networking example	35
Remembering redundancy	35
Saving the research	36
See also	36
Managing motherboards	36
Getting ready	37
How to do it...	37
How it works...	38
There's more...	39
See also	39
Selecting a chassis	40
Getting ready	40
How to do it...	40
How it works...	41
There's more...	42
Saddling up to a SAN	42
Getting ready	43
How to do it...	43
How it works...	44
There's more...	45
See also	45
Tallying up	46
Getting ready	46
How to do it...	46
How it works...	46
There's more...	47

Protecting your eggs	48
Getting ready	48
How to do it...	48
How it works...	48
There's more...	49
Chapter 2: Handling and Avoiding Downtime	51
Introduction	51
Determining acceptable losses	52
Getting ready	53
How to do it...	53
How it works...	54
Configuration – getting it right the first time	54
Getting ready	55
How to do it...	55
How it works...	56
There's more...	59
See also	59
Configuration – managing scary settings	59
Getting ready	59
How to do it...	60
How it works...	60
There's more...	61
Distinct settings	62
More information	62
See also	62
Identifying important tables	63
Getting ready	63
How to do it...	64
How it works...	65
There's more...	66
Reset stats	66
Using pgstattuple	67
See also	67
Defusing cache poisoning	68
Getting ready	68
How to do it...	69
How it works...	70
See also	72
Exploring the magic of virtual IPs	72

Getting ready	73
How to do it...	73
How it works...	74
There's more...	74
Terminating rogue connections	75
Getting ready	76
How to do it...	76
How it works...	77
There's more...	78
Reducing contention with concurrent indexes	79
Getting ready	79
How to do it...	80
How it works...	80
There's more...	80
No transactions	81
One at a time	81
Danger with OLTP use	81
See also	81
Managing system migrations	82
Getting ready	82
How to do it...	83
How it works...	84
There's more...	85
See also	85
Managing software upgrades	85
Getting ready	86
How to do it...	86
How it works...	88
There's more...	89
See also	89
Mitigating the impact of hardware failure	89
Getting ready	90
How to do it...	91
How it works...	92
There's more...	92
Copying WAL files more easily	93
Adding compression	93
Secondary delay	94
See also	94
Applying bonus kernel tweaks	94

Getting ready	95
How to do it...	95
How it works...	96
There's more...	97
Chapter 3: Pooling Resources	99
Introduction	99
Determining connection costs and limits	101
Getting ready	101
How to do it...	102
How it works...	102
There's more...	103
Installing PgBouncer	103
Getting ready	104
How to do it...	104
How it works...	106
There's more...	106
See also	107
Configuring PgBouncer safely	107
Getting ready	107
How to do it...	108
How it works...	109
There's more...	110
What about pool_mode?	110
Problems with prepared statements	111
See also	111
Connecting to PgBouncer	111
Getting ready	112
How to do it...	112
How it works...	112
See also	112
Listing PgBouncer server connections	113
Getting ready	113
How to do it...	113
How it works...	113
There's more...	115
See also	115
Listing PgBouncer client connections	115
Getting ready	116
How to do it...	116

How it works...	116
There's more...	117
See also	118
Evaluating PgBouncer pool health	118
Getting ready	118
How to do it...	118
How it works...	119
There's more...	121
See also	121
Installing pgpool	122
Getting ready	122
How to do it...	122
How it works...	124
There's more...	125
See also	125
Configuring pgpool for master/slave mode	125
Getting ready	126
How to do it...	126
How it works...	128
There's more...	129
See also	129
Testing a write query on pgpool	130
Getting ready	130
How to do it...	130
How it works...	131
There's more...	132
Swapping active nodes with pgpool	132
Getting ready	133
How to do it...	133
How it works...	134
There's more...	135
See also	136
Combining the power of PgBouncer and pgpool	136
Getting ready	136
How to do it...	137
How it works...	137
There's more...	137
Chapter 4: Troubleshooting	139
Introduction	139

Performing triage	140
Getting ready	140
How to do it...	141
How it works...	141
There's more...	143
Installing common statistics packages	143
How to do it...	144
How it works...	144
Evaluating the current disk performance with iostat	144
Getting ready	145
How to do it...	145
How it works...	145
There's more...	147
See also	147
Tracking I/O-heavy processes with iotop	147
Getting ready	148
How to do it...	148
How it works...	148
There's more...	149
See also	150
Viewing past performance with sar	150
Getting ready	150
How to do it...	150
How it works...	151
There's more...	152
See also	153
Correlating performance with dstat	153
Getting ready	153
How to do it...	153
How it works...	154
See also	156
Interpreting /proc/meminfo	156
Getting ready	157
How to do it...	157
How it works...	157
There's more...	159
See also	159
Examining /proc/net/bonding/bond0	160
Getting ready	160

How to do it...	161
How it works...	161
See also	162
Checking the pg_stat_activity view	162
Getting ready	163
How to do it...	163
How it works...	164
There's more...	166
See also	167
Checking the pg_stat_statements view	167
Getting ready	167
How to do it...	168
How it works...	169
There's more...	170
Resetting the stats	170
Catching more queries	171
See also	171
Deciphering database locks	171
Getting ready	172
How to do it...	172
How it works...	173
There's more...	173
See also	175
Debugging with strace	175
Getting ready	175
How to do it...	175
How it works...	176
There's more...	178
See also	178
Logging checkpoints properly	178
Getting ready	179
How to do it...	180
How it works...	180
There's more...	182
See also	182
Chapter 5: Monitoring	183
Introduction	183
Figuring out what to monitor	184
Getting ready	184