

Extracted from:

High Performance PostgreSQL for Rails

Reliable, Scalable, Maintainable Database Applications

This PDF file contains pages extracted from *High Performance PostgreSQL for Rails*, published by the Pragmatic Bookshelf. For more information or to purchase a paperback or PDF copy, please visit <http://www.pragprog.com>.

Note: This extract contains some colored text (particularly in code listing). This is available only in online versions of the books. The printed versions are black and white. Pagination might vary between the online and printed versions; the content is otherwise identical.

Copyright © 2023 The Pragmatic Programmers, LLC.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher.


The Pragmatic Bookshelf

Dallas, Texas

The
Pragmatic
Programmers

High Performance PostgreSQL for Rails

Reliable, Scalable, Maintainable
Database Applications



Andrew Atkinson
edited by Don N. Hagist

High Performance PostgreSQL for Rails

Reliable, Scalable, Maintainable Database Applications

Andrew Atkinson

The Pragmatic Bookshelf

Dallas, Texas



Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and The Pragmatic Programmers, LLC was aware of a trademark claim, the designations have been printed in initial capital letters or in all capitals. The Pragmatic Starter Kit, The Pragmatic Programmer, Pragmatic Programming, Pragmatic Bookshelf, PragProg and the linking *g* device are trademarks of The Pragmatic Programmers, LLC.

Every precaution was taken in the preparation of this book. However, the publisher assumes no responsibility for errors or omissions, or for damages that may result from the use of information (including program listings) contained herein.

For our complete catalog of hands-on, practical, and Pragmatic content for software developers, please visit <https://pragprog.com>.

For sales, volume licensing, and support, please contact support@pragprog.com.

For international rights, please contact rights@pragprog.com.

Copyright © 2023 The Pragmatic Programmers, LLC.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher.

ISBN-13: 979-8-88865-038-7

Encoded using the finest acid-free high-entropy binary digits.

Book version: B1.0—August 30, 2023

Preface

If you're looking to improve your skills with PostgreSQL for Ruby on Rails you've come to the right place. In this book you'll work on exercises from practical problems that will help you grow your career.

Maybe your application has grown in popularity and you've got performance problems. Maybe your database has ballooned in size and you aren't sure how to manage the growth. In the market, the traditional Database Administrator role is declining in popularity while there are more database-backed web applications than ever. Organizations from startups to huge companies choose PostgreSQL as a mission critical database to build their businesses on. Operator challenges, especially under high growth and high query volume, can put a lot of pressure on the engineering team. Why not be the team member who takes on these challenges?

Get ready to develop useful skills for operating PostgreSQL that you can immediately put to use. You'll use the latest versions of PostgreSQL and Ruby on Rails. You'll work on making your applications and databases faster, more reliable, and more resilient. As you work towards mastery of high performance operations, you'll raise the bar for operational excellence for your PostgreSQL database and your engineering team.

As you go through the book's exercises, you'll work with a real Rails application and database schema that you can see and continue to evolve as a test lab for your skill development. You'll populate millions of rows of data to work with, to help simulate your production workload. In addition to the core functionality of PostgreSQL and Ruby on Rails, you'll be working with 50 Ruby gems and PostgreSQL extensions.

Who Is This Book For?

This book is for developers looking to improve their skills with PostgreSQL and Ruby on Rails. Whether you're working on consumer scale Internet applications or enterprise B2B Software as a Service (SaaS), scaling PostgreSQL

and Rails application codebases is mission critical for the success of your business. Your team members expect you to both build with and operate these technologies, managing huge data growth, shifting business priorities, and operating reliable and cost efficient infrastructure.

If your job responsibilities or career aspirations include any of the following descriptions, you've come to the right place:

- Ruby on Rails Application Developers improving their PostgreSQL knowledge and skills
- PostgreSQL Database Administrators (DBAs) learning about Ruby on Rails and Active Record
- Infrastructure Engineers developing SQL knowledge and PostgreSQL knowledge for Replication, Parameters, Maintenance, and Partitioning
- Database Reliability Engineers (DBRE) implementing Sharding, Replication, Partitioning, and Parameter Optimization
- SQL developers experienced with other database engines looking to learn PostgreSQL
- Database engineers working with MySQL looking to learn PostgreSQL
- Web Application Developers using other full-stack web frameworks like Laravel¹ or Django,² looking to learn Ruby on Rails and Active Record
- Data Engineers working with OLAP databases interested in Replication and Change Data Capture (CDC) with PostgreSQL and OLTP use cases
- Engineers looking to get a promotion to senior levels with improved database skills

Here you'll work exclusively with the open source community distribution of PostgreSQL. Sticking with locally installed community distribution of PostgreSQL allows readers to gain confidence by developing and testing locally. After getting in some practice, readers can take their skills into production.

The focus is on Transactional workloads, also called Online Transaction Processing or (OLTP), and not Analytical workloads. OLTP workloads are the types of queries you'd see in a web application that is user facing.

What's Not Covered In This Book?

This book will help you a lot but it doesn't cover everything. It is not an introduction to PostgreSQL or Ruby on Rails. Books like *PostgreSQL: Up and Running*³ are intended to help get readers started.

1. <https://laravel.com>

2. <https://www.djangoproject.com>

3. <https://www.oreilly.com/library/view/postgresql-up-and/9781449326326/>

Pragmatic Programmers has lots of excellent books on Ruby on Rails. *Agile Web Development with Rails 7*⁴ would be a great book to go through first if you're new to Ruby on Rails. This book will help you build a good foundation with Ruby on Rails and Active Record. For readers that wish to strengthen their skills with the Ruby programming language, considering purchasing and reading *Programming Ruby 3.2 (5th Edition)*.⁵

Internals of PostgreSQL are not covered beyond some basics as needed for a chapter. Multiversion Concurrency Control (MVCC) and Isolation Levels for example are introduced with some basic information and readers are direct to the PostgreSQL documentation for more information.

User Administration, Roles, and Privileges are not covered. The GRANT command and Security concepts like Row Level Security (RLS) and Policies are not covered.

Readers will set up multiple PostgreSQL instances, however High Availability (HA) concepts like replication with Availability Zones (AZ) and Regions, or Disaster Recovery (DR) implementations are outside the scope of this book.

For concepts like Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO), consider the book *Database Reliability Engineering*⁶ which focuses on building reliable database systems.

Database Backups are critical. Your organization must collect Snapshots automatically and be able to restore a database server instance to a previous backup if disaster strikes. However, backups and restores are also outside the scope of this book beyond some basics.

In the broader PostgreSQL ecosystem, there are many options beyond what is supported in the open source community distribution where you'll run PostgreSQL on your local server instance. There are PostgreSQL forks, PostgreSQL compatible databases, and PostgreSQL extensions that make very significant modifications to how PostgreSQL operates. For example, the *Citus 11.1 open source release*⁷ modifies PostgreSQL as an extension. Although extensions like Citus do offer portions of their offerings as open source software, in general forks and different databases engines with compatibility are outside the scope of this book.

4. <https://pragprog.com/titles/rails7/agile-web-development-with-rails-7/>

5. <https://pragprog.com/titles/ruby5/programming-ruby-3-2-5th-edition/>

6. <https://www.oreilly.com/library/view/database-reliability-engineering/9781491925935/>

7. <https://www.postgresql.org/about/news/announcing-citus-11-1-open-source-release-2511/>

Ruby on Rails Skills Are In Demand

In the piece called “Big Transitions in the Tech Industry” covering Hired’s *2023 State of Software Engineers Survey*,⁸ Ruby on Rails was the most sought-after skill.

Ruby on Rails surfaced as the most in-demand skill for software engineering roles, creating 1.64x more interview requests for the developers proficient in it.

Ruby on Rails is a great way to build PostgreSQL database applications. Ruby on Rails has a Rails Guides page that is dedicated to showing how to use Active Record with PostgreSQL.⁹ Some of the capabilities are exclusive to PostgreSQL.

In the earlier days of Ruby on Rails and the Active Record ORM there was an emphasis on portability, where web applications might migrate from one database management system to another. Portability has been deemphasized over the years. Community blog posts and conference talks regularly feature capabilities that are exclusive to PostgreSQL.

Rails continues to add support for PostgreSQL capabilities to each new release. Generated Columns from PostgreSQL 12 were added as Virtual Columns¹⁰ to Active Record.

In Rails 7.1 support for Common Table Expressions (CTE) queries were added to Active Record.

These are exclusive or advanced database features that weren’t supported in the past, but have been added to the core framework. These are the kinds of topics you’ll find as you read on.

With native support for Multiple Databases, Ruby on Rails and Active Record can be used with multiple PostgreSQL databases working together. Multiple databases can be paired up using physical replication for Read and Write splitting, or multiple writeable databases can be used from the application for application level sharding. You’ll see and configure both of these.

We know skills with Ruby on Rails are in demand. What about PostgreSQL?

PostgreSQL Is A Popular Award Winner

You’ve made a great choice to invest your time learning PostgreSQL.

8. <https://hired.com/state-of-software-engineers/2023/>

9. https://guides.rubyonrails.org/active_record_postgresql.html

10. <https://blog.saeloun.com/2022/01/25/rails-7-postgres-support-for-generated-columns.html>

For deployed Ruby on Rails applications, PostgreSQL is the #1 most popular database according to the 2022 Ruby on Rails Community Survey with data from 2009 through 2022.¹¹

Besides the Rails Hosting Survey, the *DB-Engines Ranking*¹² ranks the most popular databases in the world. PostgreSQL has been a three-time #1 Winner there in 2017, 2018, and 2020.

In 2021, PostgreSQL was the second most popular database, being a runner-up to Snowflake.¹³

In 2022, PostgreSQL was the winner of the third most popular database. Snowflake and Google BigQuery¹⁴ were #1 and #2 but neither are used for OLTP.

The 2022 Stack Overflow Developer Survey¹⁵ gathered input from nearly 50,000 Professional Developers. When the developers were asked what database they used most, PostgreSQL took the #1 spot.

Your investment in yourself building skills with PostgreSQL and Ruby on Rails is a smart career move. These are very popular technologies that are in high demand, used by small startups to giant corporations, being continually improved year after year.

The future is very bright for PostgreSQL and Ruby on Rails!

11. <https://rails-hosting.com/2022/#databases>

12. <https://db-engines.com/en/ranking>

13. <https://www.snowflake.com>

14. <https://cloud.google.com/bigquery>

15. <https://survey.stackoverflow.co/2022/#most-popular-technologies-database-prof>