

PostgreSQL 15 Cookbook

*Efficient data management with
100+ recipes*

Mohammad Samsad Hussain



www.bpbonline.com

First Edition 2024

Copyright © BPB Publications, India

ISBN: 978-93-55516-61-9

All Rights Reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means or stored in a database or retrieval system, without the prior written permission of the publisher with the exception to the program listings which may be entered, stored and executed in a computer system, but they can not be reproduced by the means of publication, photocopy, recording, or by any electronic and mechanical means.

LIMITS OF LIABILITY AND DISCLAIMER OF WARRANTY

The information contained in this book is true to correct and the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but publisher cannot be held responsible for any loss or damage arising from any information in this book.

All trademarks referred to in the book are acknowledged as properties of their respective owners but BPB Publications cannot guarantee the accuracy of this information.

To View Complete
BPB Publications Catalogue
Scan the QR Code:



www.bpbonline.com

Kup książkę

Dedicated to

The readers

About the Author

Mohammad Samsad Hussain is based in Malaysia. He is a Database Architect and consultant with more than 13 years of experience in different roles (Data Analyst, Database Solution Architect, Application Consultant) among software industries. He is an IBM DB2 DBA Certified Professional and PostgreSQL, MySQL & MariaDB database Consultant as he has worked on different technologies and architectures, such as SAP HANA, Oracle, MongoDB, MS SQL and IBMi (As400).

Samsad is currently employed as Database Infrastructure Solution Architect in a Utilities based Organization in Malaysia, and he is in charge of databases and systems management and infrastructure.

About the Reviewers

- ❖ **Sami Lekesiz**, born in Turkiye, graduated from Istanbul Technical University. He has been working exclusively with PostgreSQL for the last five years of his career. His expertise relates to various PostgreSQL projects where he has played a crucial role in numerous critical tasks such as migrations, building highly available cluster architectures, and system optimizations. With his extensive knowledge and hands-on experience, Sami has been instrumental in the successful delivery of multiple PostgreSQL projects. He has always focused on delivering excellence in tasks that are central to these projects. He has ensured seamless migrations, set up highly available clusters from scratch, and implemented critical system enhancements. Looking ahead in his career, Sami aims to continue his dedicated focus on PostgreSQL, aspiring to further specialize and contribute within this domain.
- ❖ **Regina Obe** is the president of Paragon Corporation, a PostgreSQL database consulting company located in Boston, Massachusetts, USA. She has over 25 years of experience in various programming languages and database systems, especially PostgreSQL and spatial databases. She is a member of the project steering committee and development teams of two popular PostgreSQL spatial extensions, PostGIS and pgRouting. She is also a charter member and System Administration support team member for Open Source Geospatial Foundation (OSGeo), which focuses on developing and evangelizing free and open-source geospatial software. She holds a BS degree in mechanical engineering from the Massachusetts Institute of Technology.

Acknowledgement

I want to thank everyone who helped bring this book to life. Special thanks to my friend and colleague, **Abhishek Chaturvedi** and **Mohd Zaidi b. Mohd Yusoff** for their valuable contributions and strong support. Thanks to my family and friends for always cheering me on. A big thank you to the readers.

I also want to give a special thanks to the publisher of this book, BPB Publications, for believing in this project and providing the platform to make it happen. Your expertise and dedication made this dream come true. Thanks to my colleagues too, for their support and insights, adding depth to this work. Together, we have created something special, and I am truly thankful for this collaborative journey.

Preface

Welcome to **PostgreSQL 15 Cookbook**, PostgreSQL 15 stands out as a robust and reliable SQL database server. This book is here to help you understand and master PostgreSQL 15 in a practical and straightforward way.

In this book, we cover various aspects of PostgreSQL 15, providing insights and hands-on guidance. Whether you are just starting with PostgreSQL 15 or delving into advanced topics like database hierarchy, cloud provisioning, migration, and more, each chapter is crafted to improve your understanding and skills.

PostgreSQL 15 is essential in managing data. Whether you are new to it or experienced, this recipe guide makes PostgreSQL 15 accessible and valuable. We focus on real-world applications, showing the importance of each topic in practical situations.

This book is divided into 13 chapters, each focused on a specific aspect of PostgreSQL 15. Whether you are exploring the basics, advancing your skills in managing PostgreSQL 15, or troubleshooting problems, our goal is to present information that is clear, concise, and practical. Our goal is to equip you with the knowledge to use PostgreSQL 15 effectively.

Chapter 1: Up and Running with PostgreSQL 15 - In this chapter, we will learn the basics of databases and their core architecture concepts, seeing how these elements work together. The focus is on starting with PostgreSQL, giving you simple instructions for installing and setting it up on Linux VMs. We will also cover database upgrades to both minor and major releases.

Chapter 2: Database Hierarchy - We will work on creating and using databases in PostgreSQL. We will cover best practices for user access and authorization, the hierarchy of database objects, tables inheritance, concurrent indexes, and workarounds with various database objects.

Chapter 3: Cloud Provisioning - In this chapter, we will work on setting up a PostgreSQL cloud instance and handling database connections using AWS EC2 and RDS instance. We will also explain how to perform native backup and restore procedures for both AWS EC2 and RDS instances. Additionally, we will explore managing PostgreSQL on the cloud, including aspects like connection management and replication.

Chapter 4: Migration - This chapter explores different migration options for moving SQL databases to PostgreSQL and transferring on-premises data to the AWS cloud. We will work on each migration method, including the use of tools like PgLoader. Additionally, we will explore the best practices to ensure a seamless and efficient migration process.

Chapter 5: Transaction Log - We will explore admin and application-specific logs, discover practical solutions, and gain a detailed understanding of PostgreSQL transaction logs. We will examine the crucial role of the Transaction Log in PostgreSQL replication, discussing topics such as Archive mode and WAL management.

Chapter 6: Partitioning and Sharding - We will discuss in detail the concepts of scaling up your PostgreSQL database based on your organization's limitations and operational needs. The goal is to provide practical concepts for enhancing the performance of your business-critical data, covering topics such as partitioning, sharding, and the role of Foreign Data Wrapper.

Chapter 7: Replication and High Availability - In this chapter, we will learn replication and High Availability solutions for PostgreSQL, understanding the best fit for different practical scenarios and business needs. We'll also dive into the concepts of Load balancing in PostgreSQL, offering a mix of theoretical understanding and practical solutions.

Chapter 8: Leveraging SQL - In this chapter, we will talk about how to get to know and work with data in the database. Once you understand the basics of PostgreSQL, this section will guide you in using SQL to interact with relational data stored in PostgreSQL databases. We will also cover how to organize query access using psql, explore Postgres JSON Query, and understand the Postgres CAST Operator.

Chapter 9: Server Controls and Auditing - We will explore the basic concepts of authentication and encryption in the PostgreSQL database server. We'll look into common methods to encrypt your client connections and provide practical solutions for user and group control to ensure flexible access control to your database. Additionally, we have a dedicated section on database authentication through SSL.

Chapter 10: Backup - We will explore the basics of database backup, offering practical examples and insights into various backup tools, such as pg_probackup and pgBackRest. We will also discuss planning for backups and strategies to improve backup performance.

Chapter 11: Recovery - This chapter covers the fundamentals of database recovery. We will explore the concept, discuss planning for recovery, understand crash recovery, and explore Point-in-Time Recovery (PITR) and table-level recovery.

Chapter 12: Monitoring and Diagnosis - In this chapter, we will discuss how to make your setup more reliable by adding monitoring for crucial actions in your database. We will look into different monitoring approaches for database components. Additionally, we will check out tools like Prometheus and Grafana, explore statistics, find monitoring scripts, and understand the significance of fsync.

Chapter 13: Troubleshooting - This chapter will explore various situations where a PostgreSQL database can be beneficial. We will provide practical solutions to help database administrators manage their databases effectively. The topics covered include transaction logs and checkpoints, benchmarking with pgbench, and utilizing data checksums with pg_checksums.

Code Bundle and Coloured Images

Please follow the link to download the
Code Bundle and the *Coloured Images* of the book:

<https://rebrand.ly/j6vnarp>

The code bundle for the book is also hosted on GitHub at

<https://github.com/bpbpublications/PostgreSQL-15-Cookbook>.

In case there's an update to the code, it will be updated on the existing GitHub repository.

We have code bundles from our rich catalogue of books and videos available at **<https://github.com/bpbpublications>**. Check them out!

Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

errata@bpbonline.com

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

Did you know that BPB offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.bpbonline.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at :

business@bpbonline.com for more details.

At **www.bpbonline.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 BPB books and eBooks.

Piracy

If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at **business@bpbonline.com** with a link to the material.

If you are interested in becoming an author

If there is a topic that you have expertise in, and you are interested in either writing or contributing to a book, please visit **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions. We at BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

<https://discord.bpbonline.com>



Table of Contents

1. Up and Running with PostgreSQL 15	1
Introduction	1
Structure	1
Objectives	2
Introduction to PostgreSQL 15	2
Exploring the history of PostgreSQL	3
Learning about PostgreSQL architecture	3
Learning about PostgreSQL database structure	4
Installation methodology	5
<i>Recipe 1: Working with installation from binaries</i>	<i>5</i>
<i>Recipe 2: Working with installation from source code</i>	<i>6</i>
<i>Recipe 3: Database configuration file</i>	<i>8</i>
Database server and client environment	9
<i>Recipe 4: Installing and configuring pgAdmin client tool</i>	<i>10</i>
<i>Recipe 5: Working with PostgreSQL server access solution</i>	<i>12</i>
<i>Recipe 6: Working with PostgreSQL remote access</i>	<i>14</i>
<i>Recipe 7: Discovering PostgreSQL database structural object</i>	<i>15</i>
Exploring backward compatibility	17
<i>Recipe 8: Working with upgrade on minor release</i>	<i>18</i>
<i>Recipe 9: Working with upgrade on major release</i>	<i>19</i>
Conclusion	22
2. Database Hierarchy	23
Introduction	23
Structure	23
Objectives	24
Access management	24
Database instance and object hierarchy	24
<i>Recipe 10: Adaptation with PostgreSQL Schema</i>	<i>26</i>

Advanced object-relational table concept.....	29
<i>Recipe 11: Getting insight to tables inheritance.....</i>	<i>30</i>
<i>Recipe 12: Indexing technique with PostgreSQL database</i>	<i>32</i>
<i>Recipe 13: Workaround with concurrent index</i>	<i>35</i>
Common table expression	36
<i>Recipe 14: Getting insight to SELECT in WITH queries CTE</i>	<i>36</i>
Discover database storage layout	41
<i>Recipe 15: Adapt database storage for PostgreSQL database.....</i>	<i>42</i>
<i>Recipe 16: Adapt database log directory.....</i>	<i>44</i>
<i>Recipe 17: Workaround with PostgreSQL TOAST</i>	<i>45</i>
<i>Recipe 18: Parse the database admin specific start-up logs.....</i>	<i>48</i>
<i>Recipe 19: Getting insight to PostgreSQL memory configuration</i>	<i>50</i>
Local memory area	50
Shared memory area	51
Conclusion	52
3. Cloud Provisioning.....	53
Introduction	53
Structure	53
Objectives	54
Introduction to cloud solution for PostgreSQL	54
Cloud service model.....	54
Cloud deployment option	55
Managed or self-managed options.....	55
Exploring EC2 and RDS instance.....	55
<i>Recipe 20: Manage PostgreSQL instance with AWS EC2</i>	<i>56</i>
<i>Recipe 21: Managing PostgreSQL instance with AWS RDS.....</i>	<i>65</i>
<i>Recipe 22: Native backup or restore with AWS EC2 instance.....</i>	<i>72</i>
<i>Recipe 23: Backup or restore with AWS RDS instance</i>	<i>80</i>
<i>Recipe 24: Working with replication on AWS for PostgreSQL</i>	<i>88</i>
Conclusion	94

4. Migration	95
Introduction	95
Structure	95
Objectives	96
Introduction to migration for PostgreSQL	96
Exploring migration methodology.....	96
Understanding PostgreSQL database system migration types.....	97
<i>Recipe 25: On-premise to AWS EC2 instance migration</i>	<i>98</i>
<i>Database migration roadmap</i>	<i>99</i>
<i>Recipe 26: Migrating PostgreSQL from EC2 to RDS instance on AWS</i>	<i>100</i>
Getting ready with pgloader	118
<i>Recipe 27: Setting up pgloader.....</i>	<i>119</i>
<i>Recipe 28: Getting insight to migrate MariaDB to PostgreSQL on EC2 instance</i>	<i>121</i>
<i>Recipe 29: Using db2topg for migration to PostgreSQL</i>	<i>125</i>
<i>Recipe 30: Setting up foreign data wrapper for PostgreSQL migration</i>	<i>127</i>
Conclusion	136
5. Transaction Log.....	137
Introduction	137
Structure	137
Objectives	138
Introducing transaction log	138
Understanding WAL.....	139
Exploring the internal layout of WAL segment.....	139
Exploring the internal layout of XLOG record	140
Configuring the WAL file.....	141
<i>PostgreSQL WAL benefits.....</i>	<i>142</i>
<i>Recipe 31: How to enable/disable archive mode</i>	<i>143</i>
<i>Recipe 32: Working with remote WAL Archive options</i>	<i>146</i>
<i>Recipe 33: Working with WAL compression option for space management</i>	<i>148</i>
Explore SQL statement transaction sizing.....	150
<i>Recipe 34: Configuring and managing WAL performance parameter.....</i>	<i>151</i>
<i>Recipe 35: Administering continuous archiving</i>	<i>152</i>

Managing the WAL writer process.....	153
<i>PostgreSQL 15 WAL writer process versus checkpoint processing</i>	<i>154</i>
Optimizing checkpoint processing.....	154
Automating vacuum process	155
<i>Recipe 36: Getting insight to vacuum process in PostgreSQL database.....</i>	<i>157</i>
<i>Recipe 37: Debug PostgreSQL autovacuum problem</i>	<i>160</i>
Implementing continuous archiving and archive logs	161
<i>Recipe 38: PostgreSQL 15 continuous archiving and archive logs examples.....</i>	<i>161</i>
Enhancing performance and benefits of WAL.....	164
<i>Recipe 39: Proactive solution to delete PostgreSQL archive logs</i>	<i>165</i>
Conclusion	168
6. Partitioning and Sharding	171
Introduction	171
Structure	171
Objectives	172
Partitioning and sharding.....	172
Partitioning classifications.....	172
<i>Recipe 40: Setup and exploit partitioning in PostgreSQL.....</i>	<i>173</i>
<i>Recipe 41: Partition management with PostgreSQL 15</i>	<i>179</i>
<i>Recipe 42: Getting insight to vertical and horizontal partitioning.....</i>	<i>183</i>
Horizontal partitioning.....	183
Vertical partitioning.....	186
<i>Recipe 43: Working with automatic partition with PostgreSQL.....</i>	<i>187</i>
Manage partition using pg_partman extension	188
<i>Recipe 44: Manage partition using custom scripts with a scheduler</i>	<i>194</i>
<i>Recipe 45: Getting insight to partition pruning</i>	<i>197</i>
Partitioning versus sharding	200
Shard strategies and rebalancer	201
<i>Recipe 46: Configure sharding with Citus Data</i>	<i>202</i>
Rebalancer	204
Foreign data wrapper	204
Conclusion	205

7. Replication and High Availability	207
Introduction	207
Structure	207
Objectives	208
Replication and high availability	208
Understanding the CAP theorem	209
Replication classification	210
<i>Recipe 47: Scaling Postgres with primary-standby replication</i>	<i>211</i>
Using replication slots	213
<i>Recipe 48: Scaling Postgres with multi-master replication</i>	<i>214</i>
<i>Recipe 49: Setup delay standby in PostgreSQL</i>	<i>222</i>
<i>Recipe 50: Performing PITR recovery using delay standby</i>	<i>225</i>
<i>Recipe 51: PostgreSQL 15 promote standby database to primary</i>	<i>227</i>
<i>Recipe 52: Installing and configuring repmgr for PostgreSQL cluster</i>	<i>228</i>
Switchover versus failover	235
<i>Recipe 53: Switchover with repmgr for PostgreSQL</i>	<i>235</i>
<i>Recipe 54: Failover with repmgr for PostgreSQL</i>	<i>237</i>
Understanding proxy and load balancing	239
<i>Recipe 55: Deploying PostgreSQL Automatic Failover for HA solution</i>	<i>239</i>
<i>Recipe 56: Using HAProxy for HA solution</i>	<i>244</i>
Monitoring replication	248
Conclusion	248
8. Leveraging SQL	249
Introduction	249
Structure	249
Objectives	250
Exploring data access	250
Working with tables and data	251
<i>Recipe 57: Tables and Data Operations</i>	<i>251</i>
<i>Table creation</i>	<i>251</i>
<i>Data insertion</i>	<i>252</i>

Data retrieval.....	252
Data modification	253
Table operations	254
Data integrity and transactions	254
Recipe 58: Querying PostgreSQL data	255
Recipe 59: Querying using shell script in PostgreSQL	258
Exploring DML, DDL, TCL and DCL	260
Recipe 60: Database Query and Control Language	260
Data manipulation language.....	260
Data definition language	261
Transaction control language	262
Data control language	263
Recipe 61: DDL adaptation	264
Recipe 62: Working with dataset export/import	266
Recipe 63: Dataset load from spreadsheet/flat files.....	269
Recipe 64: Structuring query access with psql	270
Recipe 65: Exploring PostgreSQL join and subqueries	272
Recipe 66: Querying JSON data	274
Recipe 67: Working with PostgreSQL CAST operator.....	277
Recipe 68: Working with database consistency and integrity.....	279
Unique constraints.....	279
Foreign key constraints	280
Recipe 69: Getting insight to Python and Java connection.....	282
Recipe 70: Importing BLOB data types into PostgreSQL	287
Conclusion	289
9. Server Controls and Auditing.....	291
Introduction	291
Structure	291
Objectives	292
Introduction to server control and auditing.....	292
Recipe 71: Database privileges	293

<i>Recipe 72: PostgreSQL role management and authorization</i>	298
SSL/TLS authentication.....	302
<i>Recipe 73: Setting up SSL authentication in PostgreSQL</i>	303
Exploring TDE and encryption in PostgreSQL.....	305
<i>Recipe 74: Encryption with pgcrypto in PostgreSQL</i>	306
PostgreSQL auditing.....	310
<i>Recipe 75: Installing and configuring pgaudit</i>	310
<i>Recipe 76: Using audit log with PostgreSQL trigger</i>	314
<i>Recipe 77: Working with log_statement</i>	317
LDAP authentication.....	322
<i>Recipe 78: Getting ready with LDAP authentication</i>	322
Conclusion	324
10. Backup	325
Introduction	325
Structure	325
Objectives	326
Introduction to database backup	326
Exploring SQL dump and file system level backup.....	327
<i>Recipe 79: Working with logical backup</i>	328
<i>Recipe 80: Working with physical backup</i>	329
<i>Recipe 81: Automating backup</i>	332
Exploring backup tools	335
<i>Recipe 82: Working with pg_probackup, pgBackRest tool</i>	336
<i>Recipe 83: Installing and configuring pg_probackup for PostgreSQL</i>	336
<i>Recipe 84: Installing and configuring pgBackRest for PostgreSQL</i>	343
<i>Recipe 85: Installing and configuring Barman</i>	349
<i>Recipe 86: Incremental/differential backup</i>	356
<i>Recipe 87: Working with schema level backup</i>	360
<i>Recipe 88: Monitoring backup</i>	363
Improving backup performance	364
Conclusion	365

11. Recovery	367
Introduction	367
Structure	367
Objectives	368
Introduction to database recovery	368
Planning recovery	368
<i>Recipe 89: Recovery from logical backup.....</i>	<i>369</i>
<i>Recipe 90: Recovery from physical backup.....</i>	<i>370</i>
Understanding crash recovery	371
Recovering with PITR	372
<i>Recipe 91: Point-in-time recovery.....</i>	<i>372</i>
Incremental/differential restore	376
<i>Recipe 92: Incremental/differential restore with Barman.....</i>	<i>377</i>
<i>Recipe 93: Restoring database with Barman.....</i>	<i>383</i>
Comparing dropped versus damaged table recovery.....	385
<i>Recipe 94: Working with tables recovery</i>	<i>386</i>
<i>Recipe 95: Working with schema level restore</i>	<i>387</i>
Conclusion	391
12. Monitoring and Diagnosis	393
Introduction	393
Structure	393
Objectives	394
Introducing monitoring and diagnosis.....	394
System resource monitoring.....	394
Database monitoring	395
Discovering monitoring script	395
<i>Recipe 96: Utilizing script-based monitoring</i>	<i>395</i>
Prometheus and Grafana	397
<i>Recipe 97: Managing Prometheus and Grafana.....</i>	<i>397</i>
<i>Recipe 98: Kill/terminate hung query</i>	<i>410</i>
Explore statistics.....	412
<i>Recipe 99: Working with statistic collector and analyzer</i>	<i>412</i>

Monitoring with system view	415
<i>Recipe 100: Getting insight into query monitoring</i>	416
<i>Recipe 101: Getting insight to monitor database lock</i>	417
<i>Recipe 102: Getting insight to monitor active session</i>	419
<i>Recipe 103: Working with query plan</i>	421
Vacuum and bloat consideration	422
<i>Recipe 104: Getting insight into vacuum and bloat</i>	423
Explore fsync	427
<i>Recipe 105: Getting insight into fsync</i>	429
PostgreSQL other monitoring use cases	432
<i>Recipe 106: Monitoring cache hit ratio</i>	432
<i>Recipe 107: Monitor long-running query</i>	433
Conclusion	435
13. Troubleshooting	437
Introduction	437
Structure	437
Objectives	438
Transaction log and checkpoint	438
<i>Recipe 108: Adapt log file rotation</i>	439
<i>Recipe 109: Tune transaction log and checkpoints</i>	441
Benchmarking.....	443
<i>Recipe 110: Benchmark performance with pgbench</i>	444
<i>Recipe 111: Practice with key parameters in PostgreSQL</i>	446
Data checksum	448
<i>Recipe 112: Working with pg_checksums</i>	449
Resolution steps:	449
Conclusion	451
Index	453-460

CHAPTER 1

Up and Running with PostgreSQL 15

Introduction

PostgreSQL, aka Postgres, is an advanced open-source relational database and world's fastest-growing database management system. Being open-source, PostgreSQL is released under the PostgreSQL License, a liberal open source license, similar to the BSD or MIT licenses. As of the time of writing this book, the most recent major release of PostgreSQL is version 16.

This chapter provides an introduction to PostgreSQL 15 database systems. We begin by discussing the history of PostgreSQL. We then look at the installation, configuration and starting up PostgreSQL server. Finally, we will cover a couple of common configuration options and database client configuration.

Structure

In this chapter, we will cover the following topics:

- Introduction to PostgreSQL 15
- Exploring the history of PostgreSQL
- Learning about PostgreSQL architecture
- Learning about PostgreSQL database structure

- Installation methodology
- Database server and client environment
- Exploring backward compatibility

Objectives

This chapter covers a vast range of topic that provides an understanding of database overview and architecture of core concept and how these component work together. This chapter is all about getting up and running with PostgreSQL using basic recipes and step to cover the installation and setup of PostgreSQL on Linux VMs with step-by-step instructions. If you are already well versed with PostgreSQL basics, It is advisable to glance through the chapter quickly before proceeding to the next chapter.

Introduction to PostgreSQL 15

PostgreSQL is a powerful, flexible, and scalable general-purpose database. As it is highly stable, very low effort is required to maintain this DBMS. It runs on a wide variety of operating systems flavours, including Linux, Unix and windows.

PostgreSQL 15, the latest version released in October 2022, comes with nicer and consistent functionalities. With the recent release the new enhancement is essential reasons to upgrade. Enhancement with PostgreSQL 15 offers, few of them are:

- Addition of **MERGE** command that improves Developer experience.
- support for **LZ4** and **zstd** compression to **write-ahead log (WAL)** files
- Enhancement in logging format: **jsonlog**
- Improved collation support
- Enhancement with logical replication support for using **two-phase commit (2PC)**
- Support inspection on write-ahead log files with SQL interface
- Support for **WAL** compression with **LZ4** and **Zstandard**
- enhance performance of sorts that exceed **work_mem**

Beside of its new features, PostgreSQL is **ACID (atomicity, consistency, isolation, and durability)** complaint standard and transactional. The database programmed in C and employs a monolithic design, which means that the components are entirely united and work in a systematic manner.

Exploring the history of PostgreSQL

PostgreSQL version history dates back in mid-1980s, as a follow-up to INGRES. Postgres has undergone several significant changes since that time. Postgres used the POSTQUEL query language until 1994, In mid-1996s the project was renamed to PostgreSQL to reflect its support for SQL.

Following is the timeline of release history for PostgreSQL database with its enhancements or changes. Here, we cover the last three major versions that have not yet achieved the **end Of life (EOL)**, please refer to *Table 1.1*:

Database Release	Enhancements
PostgreSQL 13	• Incremental sorting • Parallelized vacuuming of indexes • Improved performance for queries that use aggregates or partitioned tables
PostgreSQL 14	• Support in place tablespaces • Change in default password encryption to scram-sha-256 • additions to remove expired index entries
PostgreSQL 15	• Support for merge command • Log output using JSON Format • Enhancement in in-memory • Support for zstd compression

Table 1.1: PostgreSQL database release history - enhancement

Learning about PostgreSQL architecture

The PostgreSQL architecture defines the relationship between various PostgreSQL systems. The basic architecture follows the client-server model, where the key concept is that the client makes a request and initiates the connection. The server responds and provide the service.

This architecture enables a PostgreSQL system to service a wide variety of different clients that can connect locally or over the network. To each successful client, connection spawn a new child process via fork.

PostgreSQL contain many internal components, some are fundamental to PostgreSQL and some are optional. Let us take a look around a typical PostgreSQL internal and external components in the following figure:




Figure 1.1: PostgreSQL database cluster architecture

Learning about PostgreSQL database structure

PostgreSQL Cluster, refers to the collection of database all managing the shared data directory that contain one or more database, that are managed by single server instance. A cluster is generally created for you when you install PostgreSQL database. **initdb** initializes the database cluster with default locale and character set encoding, unless custom setting is specified with the required option.

After the initialization, a newly created instance has **template0**, **template1** and postgres as the built-in databases. The database is the **template** upon which the new database is built.

In PostgreSQL, the schema is the logical grouping of database objects or the collection of databases objects are maintained and organized in database. A SCHEMA is essentially a namespace, it makes it possible to have a table with the same name if it belongs to a different schema. Let get insight to PostgreSQL:

Every database objects has two parts:

schema_name.object_name
{prefix}.{suffix}

When you create a database object, it is always created within a schema. The prefix part logically distinguish your organizational data with other database object.

The syntax for creating a schema:

```
# create schema <schema_name>;
```

or

```
# create schema IF NOT EXISTS <schema_name>;
```

Note: IF NOT EXISTS clause is an optional parameter of the CREATE SCHEMA to check whether or not the SCHEMA of the same name already exist in the database before creating it. If it already exist then it will produce an error.

Installation methodology

Now that you have an overview of the PostgreSQL database, the next step is to install it. In this section we walk you through the installation methodologies with step-by-step instructions. To offer flexibility, in this book we will perform the installation through binary packages and from the source code.

In this recipe we will cover the installation of PostgreSQL in the following 2 aspects:

- working with installation from binaries
- working with installation from source code

Recipe 1: Working with installation from binaries

In this recipe, we will cover the installation of PostgreSQL from binary package. Most of the installs of PostgreSQL are on various flavours of Linux and Windows. For this recipe, we will use Red Hat version 8.4 x86_64 Linux server.

Determine the running Linux version. The following command will give us the information we need:

```
# cat /etc/os-release
NAME="Red Hat Enterprise Linux"
VERSION="8.4 (Ootpa)"
ID="rhel"
ID_LIKE="fedora"
VERSION_ID="8.4"
```

First, assuming that we have downloaded the PostgreSQL binaries from the download site. If not then you can download the latest version of PostgreSQL from their official website at <https://www.postgresql.org/download/> after select your host operating system and version and follow the installation instructions.