

Mastering PostgreSQL 18: A Comprehensive Guide

This book provides an in-depth look into PostgreSQL 18, offering a complete walkthrough from basic concepts to advanced functionalities. Perfect for database administrators, developers, and anyone looking to harness the full power of PostgreSQL.

Table of Contents

Preface

- [What Is PostgreSQL?](#)
- [A Brief History of PostgreSQL](#)
- [Conventions](#)
- [Further Information](#)
- [Bug Reporting Guidelines](#)

Tutorial

- [Getting Started](#)
- [The SQL Language](#)
- [Advanced Features](#)

The SQL Language

- [SQL Syntax](#)
- [Data Definition](#)
- [Data Manipulation](#)
- [Queries](#)
- [Data Types](#)
- [Functions and Operators](#)
- [Type Conversion](#)
- [Indexes](#)
- [Full Text Search](#)
- [Concurrency Control](#)
- [Performance Tips](#)
- [Parallel Query](#)

Server Administration

- [Installation from Binaries](#)

- [Installation from Source Code](#)
- [Server Setup and Operation](#)
- [Server Configuration](#)
- [Client Authentication](#)
- [Database Roles](#)
- [Managing Databases](#)
- [Localization](#)
- [Routine Database Maintenance Tasks](#)
- [Backup and Restore](#)
- [High Availability, Load Balancing, and Replication](#)
- [Monitoring Database Activity](#)
- [Reliability and the Write-Ahead Log](#)
- [Logical Replication](#)
- [Just-in-Time Compilation \(JIT\)](#)

Appendices

- [Legal Notice](#)
- [The Information Schema](#)
- [Extending SQL](#)
- [Triggers](#)
- [Event Triggers](#)
- [The Rule System](#)
- [Procedural Languages](#)
- [PL/pgSQL — SQL Procedural Language](#)

PostgreSQL: Documentation: 18: 1. What Is PostgreSQL?

Navigation

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

News

- May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#)
 - [PostgreSQL 18 Beta 1 Released!](#) [Read more](#)
-

Documentation Links

PostgreSQL 18

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions

- [Current](#) (17, 16, 15, 14, 13)
- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions

- [18](#)
- [devel](#)

Unsupported Versions

- [12](#)
- [11](#)
- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)
- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)

Note: This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Section: 1. What Is PostgreSQL?

[Prev](#) [Up](#) [Home](#) [Next](#)

[Prev](#) [Up](#) [Home](#) [Next](#)

1. What Is PostgreSQL?

PostgreSQL is an object-relational database management system (ORDBMS) based on [POSTGRES, Version 4.2](#), developed at the University of California at Berkeley Computer Science Department. POSTGRES pioneered many concepts that only became available in some commercial database systems much later.

PostgreSQL is an open-source descendant of this original Berkeley code. It supports a large part of the SQL standard and offers many modern features:

- [complex queries](#)
- [foreign keys](#)
- [triggers](#)
- [updatable views](#)
- [transactional integrity](#)
- [multiversion concurrency control](#)

Also, **PostgreSQL** can be extended by the user in many ways, for example by adding new:

- [data types](#)
- [functions](#)
- [operators](#)
- [aggregate functions](#)
- [index methods](#)
- [procedural languages](#)

And because of the liberal license, **PostgreSQL** can be used, modified, and distributed by anyone free of charge for any purpose, be it private,

commercial, or academic.

Navigation

[Prev](#) | [Up](#) | [Next](#)

Current section: Preface

[Home](#) | [2. A Brief History of PostgreSQL](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

© 1996-2025 The PostgreSQL Global Development Group

2. A Brief History of PostgreSQL

[Prev](#)

[Up](#)

[Next](#)

[2. A Brief History of PostgreSQL](#) [Preface](#) [Home](#)

2. A Brief History of PostgreSQL

2.1. The Berkeley POSTGRES Project

The object-relational database management system now known as *PostgreSQL* is derived from the *POSTGRES* package written at the University of California at Berkeley. With decades of development behind it, *PostgreSQL* is now the most advanced open-source database available anywhere.

Another take on the history presented here can be found in Dr. Joe Hellerstein's paper “*Looking Back at Postgres*” [hell18](#).

2.1. The Berkeley POSTGRES Project

The *POSTGRES* project, led by Professor Michael Stonebraker, was sponsored by the Defense Advanced Research Projects Agency (DARPA), the Army Research Office (ARO), the National Science Foundation (NSF), and ESL, Inc. The implementation of *POSTGRES* began in 1986. The initial concepts for the system were presented in [ston86](#), and the definition of the initial data model appeared in [rowe87](#). The design of

the rule system at that time was described in [ston87a](#). The rationale and architecture of the storage manager were detailed in [ston87b](#).

POSTGRES has undergone several major releases since then. The first “demoware” system became operational in 1987 and was shown at the 1988 ACM-SIGMOD Conference. Version 1, described in [ston90a](#), was released to a few external users in June 1989. In response to a critique of the first rule system ([ston89](#)), the rule system was redesigned ([ston90b](#)), and Version 2 was released in June 1990 with the new rule system. Version 3 appeared in 1991 and added support for multiple storage managers, an improved query executor, and a rewritten rule system. For the most part, subsequent releases until *Postgres95* (see below) focused on portability and reliability.

POSTGRES has been used to implement many different research and production applications. These include: a financial data analysis system, a jet engine performance monitoring package, an asteroid tracking database, a medical information database, and several geographic information systems. *POSTGRES* has also been used as an educational tool at several universities. Finally, Illustra Information Technologies (later merged into [IBM](#)), which is now owned by [IBM](#), picked up the code and commercialized it. In late 1992, *POSTGRES* became the primary data manager for the Sequoia 2000 scientific computing project described in [ston92](#).

The size of the external user community nearly doubled during 1993. It became increasingly obvious that maintenance of the prototype code and support was taking up large amounts of time that should have been devoted to database research. In an effort to reduce this support burden, the Berkeley *POSTGRES* project officially ended with Version 4.2.

2.2. Postgres95

In 1994, Andrew Yu and Jolly Chen added an SQL language interpreter to *POSTGRES*. Under a new name, *Postgres95* was subsequently released

to the web to find its own way in the world as an open-source descendant of the original *POSTGRES* Berkeley code.

Postgres95 code was completely ANSI C and trimmed in size by 25%. Many internal changes improved performance and maintainability. *Postgres95* release 1.0.x ran about 30–50% faster on the Wisconsin Benchmark compared to *POSTGRES*, Version 4.2. Apart from bug fixes, the following were the major enhancements:

- The query language PostQUEL was replaced with SQL (implemented in the server). (Interface library `libpq` was named after PostQUEL.) Subqueries were not supported until *PostgreSQL* (see below), but they could be imitated in *Postgres95* with user-defined SQL functions. Aggregate functions were re-implemented. Support for the `GROUP BY` query clause was also added.
 - A new program (`psql`) was provided for interactive SQL queries, which used GNU *Readline*. This largely superseded the old *monitor* program.
 - A new front-end library, `libpgtcl`, supported Tcl-based clients. A sample shell, `pgtclsh`, provided new Tcl commands to interface *Tcl* programs with the *Postgres95* server.
 - The large-object interface was overhauled. The inversion large objects were the only mechanism for storing large objects. (The inversion file system was removed.)
 - The instance-level rule system was removed. Rules were still available as rewrite rules.
 - A short tutorial introducing regular SQL features as well as those of *Postgres95* was distributed with the source code.
 - GNU make (instead of BSD make) was used for the build. Also, *Postgres95* could be compiled with an unpatched GCC (data alignment of doubles was fixed).
-

2.3. PostgreSQL

By 1996, it became clear that the name “Postgres95” would not stand the test of time. We chose a new name, *PostgreSQL*, to reflect the relationship between the original *POSTGRES* and the more recent versions with SQL capability. At the same time, we set the version numbering to start at 6.0, putting the numbers back into the sequence originally begun by the Berkeley *POSTGRES* project.

Postgres is still considered an official project name, both because of tradition and because people find it easier to pronounce *Postgres* than *PostgreSQL*.

The emphasis during development of *Postgres95* was on identifying and understanding existing problems in the server code. With *PostgreSQL*, the emphasis has shifted to augmenting features and capabilities, although work continues in all areas.

Details about what has happened in each *PostgreSQL* release since then can be found at <https://www.postgresql.org/docs/release/>.

Navigation

[Prev](#)

[Up](#)

[Next](#)

[1. What Is PostgreSQL?](#) [Home](#) [3. Conventions](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

[PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

3. Conventions

The following conventions are used in the synopsis of a command: brackets (`[` and `]`) indicate optional parts. Brace (`{` and `}`) and vertical lines (`|`) indicate that you must choose one alternative. Dots (`...`) mean that the preceding element can be repeated. All other symbols, including parentheses, should be taken literally.

Where it enhances the clarity, SQL commands are preceded by the prompt `=>`, and shell commands are preceded by the prompt `$`. Normally, prompts are not shown, though.

An *administrator* is generally a person who is in charge of installing and running the server. A *user* could be anyone who is using, or wants to use, any part of the *PostgreSQL* system. These terms should not be interpreted too narrowly; this book does not have fixed presumptions about system administration procedures.

Prev	Up	Next
2. A Brief History of PostgreSQL		Preface 4. Further Information

This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

3. Conventions	Prev	Up	Next
history.html		preface.html	resources.html

3.	Prev	Up	Next
Conventions			
	2. A Brief History of PostgreSQL	Home	4. Further Information

Note: The original content contains many nested divs, navigation, and links that are preserved here as plain Markdown. Due to the complex HTML structure, some visual formatting is simplified.

PostgreSQL: Documentation: 18: 4. Further Information

Navigation Menu

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

Latest News

- May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#)
 - [PostgreSQL 18 Beta 1 Released!](#)
-

Documentation Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)
- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions:

- [18](#)
- [devel](#)

Unsupported Versions:

- [12](#)
- [11](#)
- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)
- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)
- [8.1](#)
- [8.0](#)
- [7.4](#)
- [7.3](#)
- [7.2](#)
- [7.1](#)

Important Notice

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Further Information

[Prev](#) [Up](#)

4. Further Information

Besides the documentation, that is, this book, there are other resources about PostgreSQL:

Resources

- **Wiki**
The [PostgreSQL wiki](#) contains the project's [FAQ](#) (Frequently Asked Questions) list, [TODO](#) list, and detailed information about many more topics.
- **Web Site**
The [PostgreSQL web site](#) carries details on the latest release and other information to make your work or play with PostgreSQL more productive.
- **Mailing Lists**
The mailing lists are a good place to have your questions answered, share experiences with other users, and contact the developers. Consult the [PostgreSQL](#) web site for details.
- **Yourself!**
PostgreSQL is an open-source project. As such, it depends on the user community for ongoing support. As you begin to use PostgreSQL, you will rely on others for help, either through the documentation or through the mailing lists. Consider contributing

your knowledge back. Read the mailing lists and answer questions. If you learn something which is not in the documentation, write it up and contribute it. If you add features to the code, contribute them.

Navigation

Prev

Up

Next

[notation.html](#) [preface.html](#) [bug-reporting.html](#)

Additional Notes

3. Conventions

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

© 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: 5. Bug Reporting Guidelines

Navigation

[Prev](#)

[Up](#)

[Next](#)

4. Further Information [Home](#) Part I. Tutorial

Bug Reporting Guidelines

5.1. Identifying Bugs

When you find a bug in *PostgreSQL*, we want to hear about it. Your bug reports are essential in making *PostgreSQL* more reliable because even the utmost care cannot guarantee that every part of *PostgreSQL* will work on every platform under every circumstance.

The following suggestions are intended to help you form bug reports that can be handled effectively. While not mandatory, following them is generally advantageous.

We cannot promise to fix every bug immediately. Critical bugs affecting many users or obvious issues are prioritized. Sometimes we tell you to update to a newer version to see if the bug persists, or decide it cannot be fixed until a major rewrite is done. If you need immediate help, consider obtaining a commercial support contract.

5.1. Recognizing Bugs

Before reporting, please verify the documentation to ensure what you are trying is correct. If the documentation is unclear or a program behaves differently than described, that is a bug.

Circumstances considered bugs include, but are not limited to:

- Program terminates with a fatal signal or operating system error message pointing to a problem.
- Produces wrong output.
- Refuses valid input.
- Accepts invalid input without error.
- Fails to compile, build, or install on supported platforms.

“Program” refers to any executable, not only the backend process.

Slow or resource-consuming behavior is not necessarily a bug. Check documentation or ask on mailing lists for tuning advice. Non-compliance with SQL standards is not always a bug unless explicitly claimed.

Always review the TODO list and FAQ for known issues before proceeding.

5.2. What to Report

The key is to state facts clearly. Avoid speculation or vague descriptions.

Every bug report should include:

- **Exact steps to reproduce** from start-up, self-contained, including necessary SQL commands and data setup.
 - **Output received**, with error messages and logs if applicable.
 - **Expected output**.
 - **Command-line options** and configuration details.
 - **Differences from standard installation**.
-

- **PostgreSQL version** (via `SELECT version();` or `--version`).
- **Platform info** (kernel, libraries, hardware).

Be thorough: lengthy reports are better than missing crucial info.
Consider consulting [this article](#) for more tips.

5.3. Where to Report Bugs

- Send reports to the mailing list: `<pgsql-bugs@lists.postgresql.org>`.
- Use the web form at [PostgreSQL bug report page](#).
- For security issues, report privately to `<security@postgresql.org>`.
- Do **not** send bug reports to user lists like `<pgsql-sql@lists.postgresql.org>` or `<pgsql-general@lists.postgresql.org>`.
- Do **not** send reports to `<pgsql-hackers@lists.postgresql.org>` unless needed for further review.
- Documentation bugs should be reported to `<pgsql-docs@lists.postgresql.org>`.
- Porting issues go to `<pgsql-hackers@lists.postgresql.org>`.

Note: Due to spam, all mailing lists are moderated, so there may be delays. To subscribe, visit <https://lists.postgresql.org/>.

End of Page

PostgreSQL: Documentation: 18: Chapter 1. Getting Started

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

[Current](#) (17) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

[12](#) / [11](#) / [10](#) / [9.6](#) / [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#) / [7.4](#) / [7.3](#) / [7.2](#)

*This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the
other supported versions listed above instead.*

Chapter 1. Getting Started

Table of Contents

- [1.1. Installation](#)
 - [1.2. Architectural Fundamentals](#)
 - [1.3. Creating a Database](#)
 - [1.4. Accessing a Database](#)
-

[Prev](#)

[Up](#)

[Next](#)

Part I. Tutorial [Home](#) 1.1. Installation

*This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the
other supported versions listed above instead.*

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 2. The SQL Language

[Home](#)
[About](#)
[Download](#)
[Documentation](#)
[Community](#)
[Developers](#)
[Support](#)
[Donate](#)
[Your account](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation → [PostgreSQL 18](#)

Supported Versions

[Current](#) ([17](#))

/

[16](#)

/

[15](#)

/

[14](#)

/

[13](#)

Development Versions

[18](#)

/

[devel](#)

Unsupported versions

[12](#)

/

[11](#)

/

[10](#)

/

[9.6](#)

/

[9.5](#)

/

[9.4](#)

/

[9.3](#)

/

[9.2](#)

/

[9.1](#)

/

[9.0](#)

/

[8.4](#)

/

[8.3](#)

/

[8.2](#)

/

[8.1](#)

/

[8.0](#)

/

[7.4](#)

/

[7.3](#)

/

[7.2](#)

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of
the other supported versions listed above instead.

Chapter 2. The SQL Language

Table of Contents

- [2.1. Introduction](#)
- [2.2. Concepts](#)
- [2.3. Creating a New Table](#)
- [2.4. Populating a Table With Rows](#)
- [2.5. Querying a Table](#)
- [2.6. Joins Between Tables](#)
- [2.7. Aggregate Functions](#)
- [2.8. Updates](#)
- [2.9. Deletions](#)

Footer

PostgreSQL:

Documentation: 18:

Chapter 3. Advanced Features

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)
- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions:

- [18](#)
- [devel](#)

Unsupported versions:

- [12](#)
- [11](#)
- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)
- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)
- [8.1](#)
- [8.0](#)
- [7.4](#)
- [7.3](#)
- [7.2](#)

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of
the other supported versions listed above instead.

Chapter 3. Advanced Features

[Prev](#)

[Up](#)

Chapter 3. Advanced Features

Table of Contents

- [3.1. Introduction](#)
- [3.2. Views](#)
- [3.3. Foreign Keys](#)
- [3.4. Transactions](#)
- [3.5. Window Functions](#)
- [3.6. Inheritance](#)
- [3.7. Conclusion](#)

[Prev](#)

[Up](#)

[Next](#)

[2.9. Deletions](#) [Home](#) [3.1. Introduction](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 4. SQL Syntax

[Home](#) | [About PostgreSQL](#) | [Contact](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)
- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions:

- [18](#)
- [devel](#)

Unsupported versions:

- [12](#)
- [11](#)
- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)

- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)
- [8.1](#)
- [8.0](#)
- [7.4](#)
- [7.3](#)
- [7.2](#)
- [7.1](#)

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of
the other supported versions listed above instead.

Chapter 4. SQL Syntax

Prev Up Next

[Prev](#) [Up](#) [Next](#)

Chapter 4. SQL Syntax

Table of Contents

- [4.1. Lexical Structure](#)
 - [4.1.1. Identifiers and Key Words](#)
 - [4.1.2. Constants](#)
 - [4.1.3. Operators](#)
 - [4.1.4. Special Characters](#)
 - [4.1.5. Comments](#)
 - [4.1.6. Operator Precedence](#)
- [4.2. Value Expressions](#)

- 4.2.1. Column References
 - 4.2.2. Positional Parameters
 - 4.2.3. Subscripts
 - 4.2.4. Field Selection
 - 4.2.5. Operator Invocations
 - 4.2.6. Function Calls
 - 4.2.7. Aggregate Expressions
 - 4.2.8. Window Function Calls
 - 4.2.9. Type Casts
 - 4.2.10. Collation Expressions
 - 4.2.11. Scalar Subqueries
 - 4.2.12. Array Constructors
 - 4.2.13. Row Constructors
 - 4.2.14. Expression Evaluation Rules
 - 4.3. Calling Functions
 - 4.3.1. Using Positional Notation
 - 4.3.2. Using Named Notation
 - 4.3.3. Using Mixed Notation
-

This chapter describes the syntax of SQL.

It forms the foundation for understanding the following chapters which will go into detail about how SQL commands are applied to define and modify data.

We also advise users who are already familiar with SQL to read this chapter carefully because it contains several rules and concepts that are implemented inconsistently among SQL databases or that are specific to *PostgreSQL*.

Support Note

This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of

the other supported versions listed above instead.

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group



Chapter 5. Data Definition

Table of Contents

1. [5.1. Table Basics](#)
2. [5.2. Default Values](#)
3. [5.3. Identity Columns](#)
4. [5.4. Generated Columns](#)
5. [5.5. Constraints](#)
 - [5.5.1. Check Constraints](#)
 - [5.5.2. Not-Null Constraints](#)
 - [5.5.3. Unique Constraints](#)
 - [5.5.4. Primary Keys](#)
 - [5.5.5. Foreign Keys](#)
 - [5.5.6. Exclusion Constraints](#)
6. [5.6. System Columns](#)
7. [5.7. Modifying Tables](#)
 - [5.7.1. Adding a Column](#)
 - [5.7.2. Removing a Column](#)
 - [5.7.3. Adding a Constraint](#)
 - [5.7.4. Removing a Constraint](#)
 - [5.7.5. Changing a Column's Default Value](#)
 - [5.7.6. Changing a Column's Data Type](#)
 - [5.7.7. Renaming a Column](#)
 - [5.7.8. Renaming a Table](#)
8. [5.8. Privileges](#)
9. [5.9. Row Security Policies](#)
10. [5.10. Schemas](#)

- [5.10.1. Creating a Schema](#)
- [5.10.2. The Public Schema](#)
- [5.10.3. The Schema Search Path](#)
- [5.10.4. Schemas and Privileges](#)
- [5.10.5. The System Catalog Schema](#)
- [5.10.6. Usage Patterns](#)
- [5.10.7. Portability](#)
- 11. [5.11. Inheritance](#)
 - [5.11.1. Caveats](#)
- 12. [5.12. Table Partitioning](#)
 - [5.12.1. Overview](#)
 - [5.12.2. Declarative Partitioning](#)
 - [5.12.3. Partitioning Using Inheritance](#)
 - [5.12.4. Partition Pruning](#)
 - [5.12.5. Partitioning and Constraint Exclusion](#)
 - [5.12.6. Best Practices for Declarative Partitioning](#)
- 13. [5.13. Foreign Data](#)
- 14. [5.14. Other Database Objects](#)
- 15. [5.15. Dependency Tracking](#)

This chapter covers how one creates the database structures that will hold one's data. In a relational database, the raw data is stored in tables, so the majority of this chapter is devoted to explaining how tables are created and modified and what features are available to control what data is stored in the tables. Subsequently, we discuss how tables can be organized into schemas, and how privileges can be assigned to tables. Finally, we will briefly look at other features that affect the data storage, such as inheritance, table partitioning, views, functions, and triggers.

Prev	Calling Functions	Up	Table Basics	Next	Basic Table Definitions
4.3. Calling Functions					

Note: This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 14. Performance Tips

Navigation Menu

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

Latest News

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

[Current \(17\)](#) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

[12](#) / [11](#) / [10](#) / [9.6](#) / [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#) / [7.4](#)
/ [7.3](#) / [7.2](#) / [7.1](#)

Note

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 14. Performance Tips

Navigation Buttons

[Prev](#)

[Up](#)

[Next](#)

Chapter 14. Performance Tips

Table of Contents

- 14.1. Using [EXPLAIN](#)
-

- [14.1.1. EXPLAIN Basics](#)
 - [14.1.2. EXPLAIN ANALYZE](#)
 - [14.1.3. Caveats](#)
 - **14.2. Statistics Used by the Planner**
 - [14.2.1. Single-Column Statistics](#)
 - [14.2.2. Extended Statistics](#)
 - **14.3. Controlling the Planner with Explicit `JOIN` Clauses**
[explicit-joins.html](#)
 - **14.4. Populating a Database**
 - [14.4.1. Disable Autocommit](#)
 - [14.4.2. Use `COPY`](#)
 - [14.4.3. Remove Indexes](#)
 - [14.4.4. Remove Foreign Key Constraints](#)
 - [14.4.5. Increase `maintenance\_work\_mem`](#)
 - [14.4.6. Increase `max\_wal\_size`](#)
 - [14.4.7. Disable WAL Archival and Streaming Replication](#)
 - [14.4.8. Run `ANALYZE` Afterwards](#)
 - [14.4.9. Some Notes about `pg\_dump`](#)
 - **14.5. Non-Durable Settings**
[non-durability.html](#)
-

Performance Tips Details

Query performance can be affected by many things. Some of these can be controlled by the user, while others are fundamental to the underlying design of the system. This chapter provides some hints about understanding and tuning *PostgreSQL* performance.

Navigation

[Prev](#) | [Up](#) | [Next](#)

13.7. Locking and Indexes

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 15. Parallel Query

Navigation

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

Latest News:

[PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Breadcrumbs

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions

[Current](#) (17), 16, 15, 14, 13

Development Versions

[18](#), [devel](#)

Unsupported Versions

[12](#), [11](#), [10](#), [9.6](#)

Important Notice

This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 15. Parallel Query

[Prev](#)

[Home](#)

[Next](#)

[non-durability.html](#) [index.html](#) [how-parallel-query-works.html](#)

Chapter 15. Parallel Query Content

Table of Contents

- [15.1. How Parallel Query Works](#)
 - [15.2. When Can Parallel Query Be Used?](#)
 - [15.3. Parallel Plans](#)
 - [15.3.1. Parallel Scans](#)
 - [15.3.2. Parallel Joins](#)
 - [15.3.3. Parallel Aggregation](#)
 - [15.3.4. Parallel Append](#)
 - [15.3.5. Parallel Plan Tips](#)
 - [15.4. Parallel Safety](#)
 - [15.4.1. Parallel Labeling for Functions and Aggregates](#)
-

Section: 15.1 — How Parallel Query Works

PostgreSQL can devise query plans that leverage multiple CPUs to answer queries faster. This feature is known as *parallel query*. Many queries cannot benefit from parallel query, either due to limitations of the current implementation or because there is no faster query plan than the serial version. However, when beneficial, the speedup can be significant—sometimes more than four times faster. Queries touching large data amounts but returning few rows to the user tend to benefit most.

This chapter explains how parallel query works and when it can be used, helping users understand what to expect.

Navigation Summary

[Prev](#)

[Up](#)

[Next](#)

[non-durability.html](#) [sql.html](#) [how-parallel-query-works.html](#)

[Previous Section](#)

[Home](#)

[Next Section](#)

14.5. Non-Durable Settings [Home](#) 15.1. How Parallel Query Works

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 6. Data Manipulation

- [\[Home\]](#)(/) - [\[About\]](#)(/about/) - [\[Download\]](#)(/download/) -
[\[Documentation\]](#)(/docs/) - [\[Community\]](#)(/community/) - [\[Developers\]](#)
(/developer/) - [\[Support\]](#)(/support/) - [\[Donate\]](#)(/about/donate/) - [\[Your account\]](#)(/account/)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)
- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions:

- [18](#)
- [devel](#)

Unsupported versions:

- [12](#)
- [11](#)

- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)
- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)
- [8.1](#)
- [8.0](#)
- [7.4](#)
- [7.3](#)

This documentation is for an unsupported version of PostgreSQL.
 You may want to view the same page for the [current](#) version, or one of
 the other supported versions listed above instead.

Chapter 6. Data Manipulation	Prev	Up	Home	Next
	ddl- depend.html	sql.ht ml	index.ht ml	dml- insert.html

Chapter 6. Data Manipulation

Table of Contents

- [6.1. Inserting Data](#)
- [6.2. Updating Data](#)
- [6.3. Deleting Data](#)
- [6.4. Returning Data from Modified Rows](#)

The previous chapter discussed how to create tables and other structures to hold your data. Now it is time to fill the tables with data. This chapter covers how to insert, update, and delete table data. The chapter after this will finally explain how to extract your long-lost data from the database.

[Prev](#) | [Up](#) | [Next](#)

[Dependency Tracking](#) **[Home](#)** **[Inserting Data](#)**

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 7. Queries

Navigation

- [Home](#)
- [About](#)
- [Download](#)
- [Documentation](#)
- [Community](#)
- [Developers](#)
- [Support](#)
- [Donate](#)
- [Your account](#)

Recent News

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation for PostgreSQL 18

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)

- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions:

- [18](#)
- [devel](#)

Unsupported Versions:

- [12](#)
- [11](#)
- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)
- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)
- [8.1](#)
- [8.0](#)
- [7.4](#)
- [7.3](#)
- [7.2](#)
- [7.1](#)

Warning

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 7. Queries

Navigation

[Prev](#) [Up](#) [Next](#)

[Previous](#) [Home](#) [Next](#)

Sections

7.1 Overview

7.2 Table Expressions

- [7.2.1 The FROM Clause](#)
- [7.2.2 The WHERE Clause](#)
- [7.2.3 The GROUP BY and HAVING Clauses](#)
- [7.2.4 GROUPING SETS, CUBE, and ROLLUP](#)
- [7.2.5 Window Function Processing](#)

7.3 Select Lists

- [7.3.1 Select-List Items](#)
- [7.3.2 Column Labels](#)
- [7.3.3 DISTINCT](#)

7.4 Combining Queries

[UNION](#)

7.5 Sorting Rows

[ORDER BY](#)

7.6 **LIMIT** and **OFFSET**

[queries-limit.html](#)

7.7 **VALUES** Lists

[queries-values.html](#)

7.8 **WITH** Queries (Common Table Expressions)

[queries-with.html](#)

- 7.8.1 **SELECT** in **WITH**
- 7.8.2 Recursive Queries
- 7.8.3 CTE Materialization
- 7.8.4 Data-Modifying Statements in **WITH**

Content

The previous chapters explained how to create tables, how to fill them with data, and how to manipulate that data. Now we finally discuss how to retrieve the data from the database.

Navigation

[Prev](#)

[Home](#)

[Next](#)

[6.4. Returning Data from Modified Rows](#) [Home](#) [7.1. Overview](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

Chapter 8. Data Types

Table of Contents

- [8.1. Numeric Types](#)
 - [8.1.1. Integer Types](#)
 - [8.1.2. Arbitrary Precision Numbers](#)
 - [8.1.3. Floating-Point Types](#)
 - [8.1.4. Serial Types](#)
- [8.2. Monetary Types](#)
- [8.3. Character Types](#)
- [8.4. Binary Data Types](#)
 - [8.4.1. `bytea` Hex Format](#)
 - [8.4.2. `bytea` Escape Format](#)
- [8.5. Date/Time Types](#)
 - [8.5.1. Date/Time Input](#)
 - [8.5.2. Date/Time Output](#)
 - [8.5.3. Time Zones](#)
 - [8.5.4. Interval Input](#)
 - [8.5.5. Interval Output](#)
- [8.6. Boolean Type](#)
- [8.7. Enumerated Types](#)
 - [8.7.1. Declaration of Enumerated Types](#)
 - [8.7.2. Ordering](#)
 - [8.7.3. Type Safety](#)
 - [8.7.4. Implementation Details](#)
- [8.8. Geometric Types](#)
 - [8.8.1. Points](#)
 - [8.8.2. Lines](#)
 - [8.8.3. Line Segments](#)
 - [8.8.4. Boxes](#)
 - [8.8.5. Paths](#)

- 8.8.6. Polygons
 - 8.8.7. Circles
- 8.9. Network Address Types
 - 8.9.1. `inet`
 - 8.9.2. `cidr`
 - 8.9.3. `inet` vs. `cidr`
 - 8.9.4. `macaddr`
 - 8.9.5. `macaddr8`
- 8.10. Bit String Types
- 8.11. Text Search Types
 - 8.11.1. `tsvector`
 - 8.11.2. `tsquery`
- 8.12. UUID Type
- 8.13. XML Type
 - 8.13.1. Creating XML Values
 - 8.13.2. Encoding Handling
 - 8.13.3. Accessing XML Values
- 8.14. JSON Types
 - 8.14.1. JSON Input and Output Syntax
 - 8.14.2. Designing JSON Documents
 - 8.14.3. `jsonb` Containment and Existence
 - 8.14.4. `jsonb` Indexing
 - 8.14.5. `jsonb` Subscripting
 - 8.14.6. Transforms
 - 8.14.7. `jsonpath` Type
- 8.15. Arrays
 - 8.15.1. Declaration of Array Types
 - 8.15.2. Array Value Input
 - 8.15.3. Accessing Arrays
 - 8.15.4. Modifying Arrays
 - 8.15.5. Searching in Arrays
 - 8.15.6. Array Input and Output Syntax
- 8.16. Composite Types
 - 8.16.1. Declaration of Composite Types
 - 8.16.2. Constructing Composite Values
 - 8.16.3. Accessing Composite Types
 - 8.16.4. Modifying Composite Types
 - 8.16.5. Using Composite Types in Queries

- 8.16.6. Composite Type Input and Output Syntax
- 8.17. Range Types
 - 8.17.1. Built-in Range and Multirange Types
 - 8.17.2. Examples
 - 8.17.3. Inclusive and Exclusive Bounds
 - 8.17.4. Infinite (Unbounded) Ranges
 - 8.17.5. Range Input/Output
 - 8.17.6. Constructing Ranges and Multiranges
 - 8.17.7. Discrete Range Types
 - 8.17.8. Defining New Range Types
 - 8.17.9. Indexing
 - 8.17.10. Constraints on Ranges
- 8.18. Domain Types
- 8.19. Object Identifier Types
- 8.20. `pg_lsn` Type
- 8.21. Pseudo-Types

PostgreSQL has a rich set of native data types available to users. Users can add new types to PostgreSQL using the [CREATE TYPE](#) command.

[Table 8.1](#) shows all the built-in general-purpose data types. Most of the alternative names listed in the "Aliases" column are the names used internally by PostgreSQL for historical reasons. In addition, some internally used or deprecated types are available, but are not listed here.

Table 8.1. Data Types

Name	Aliases	Description
<code>bigint</code>	<code>int8</code>	signed eight-byte integer
<code>bigserial</code>	<code>serial8</code>	autoincrementing eight-byte integer
<code>bit [(n)</code> <code>]</code>		fixed-length bit string
<code>bit varying [(n)</code> <code>]</code>	<code>varbit [(n)</code> <code>]</code>	variable-length bit string

Name	Aliases	Description
boolean	bool	logical Boolean (true/false)
box		rectangular box on a plane
bytea		binary data (“byte array”)
character [(n)]	char [(n)]	fixed-length character string
character varying [(n)]	varchar [(n)]	variable-length character string
cidr		IPv4 or IPv6 network address
circle		circle on a plane
date		calendar date (year, month, day)
double precision	float8	double precision floating-point number (8 bytes)
inet		IPv4 or IPv6 host address
integer	int, int4	signed four-byte integer
interval [fields] [(p)]		time span
json		textual JSON data
jsonb		binary JSON data, decomposed
line		infinite line on a plane
lseg		line segment on a plane
macaddr		MAC (Media Access Control) address

Name	Aliases	Description
macaddr8		MAC (Media Access Control) address (EUI-64 format)
money		currency amount
numeric [(p), (s)]	decimal [(p), (s)]	exact numeric of selectable precision
path		geometric path on a plane
pg_lsn		PostgreSQL Log Sequence Number
pg_snapshot		user-level transaction ID snapshot
point		geometric point on a plane
polygon		closed geometric path on a plane
real	float4	single precision floating-point number (4 bytes)
smallint	int2	signed two-byte integer
smallserial	serial2	autoincrementing two-byte integer
serial	serial4	autoincrementing four-byte integer
text		variable-length character string
time [(p)] [without time zone]		time of day (no time zone)
time [(p)] with	timetz	time of day, including time zone

Name	Aliases	Description
<code>time zone</code>		
<code>timestamp [</code> <code>(p)] [</code> <code>without time zone</code> <code>]</code>		date and time (no time zone)
<code>timestamp [</code> <code>(p)] with</code> <code>time zone</code>	<code>timestamptz</code>	date and time, including time zone
<code>tsquery</code>		text search query
<code>tsvector</code>		text search document
		user-level transaction ID snapshot
<code>txid_snapshot</code>		(deprecated; see <code>pg_snapshot</code>)
<code>uuid</code>		universally unique identifier
<code>xml</code>		XML data

Compatibility

The following types (or spellings thereof) are specified by SQL:

`bigint`, `bit`, `bit varying`, `boolean`, `char`, `character varying`, `character`, `varchar`, `date`, `double precision`, `integer`, `interval`, `numeric`, `decimal`, `real`, `smallint`, `time` (with or without time zone), `timestamp` (with or without time zone), `xml`.

Each data type has an external representation determined by its input and output functions. Many of the built-in types have obvious external formats. However, several types are either unique to PostgreSQL, such as geometric paths, or have several possible formats, such as the date and time types. Some of the input and output functions are not invertible, i.e., the result of an output function might lose accuracy when compared to the original input.

[Prev](#) [Up](#) [Next](#)

7.8. [WITH](#) Queries (Common Table Expressions) [Home](#) 8.1. Numeric Types

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL:

Documentation: 18:

Chapter 9. Functions and Operators

Navigation

[Prev](#)

[Up](#)

[Next](#)

[datatype-pseudo.html](#) [sql.html](#) [functions-logical.html](#)

[Chapter 9. Functions and Operators](#) [Home](#) [Part II. The SQL Language](#)

[8.21. Pseudo-Types](#)

[9.1. Logical Operators](#)

Content

Overview

PostgreSQL provides a large number of functions and operators for the built-in data types. This chapter describes most of them, although additional special-purpose functions appear in relevant sections of the manual. Users can also define their own functions and operators, as described in [Part V](#). The `psql` commands `\df` and `\do` can be used to list all available functions and operators, respectively.

The notation used throughout this chapter to describe the argument and result data types of a function or operator is like this:

```
repeat (text, integer) → text
```

which says that the function `repeat` takes one `text` and one `integer` argument and returns a result of type `text`. The right arrow is also used to indicate the result of an example, thus:

```
repeat('Pg', 4) → PgPgPgPg
```

If you are concerned about portability then note that most of the functions and operators described in this chapter, with the exception of the most trivial arithmetic and comparison operators and some explicitly marked functions, are not specified by the SQL standard. Some of this extended functionality is present in other SQL database management systems, and in many cases this functionality is compatible and consistent between the various implementations.

Footer Navigation

[Prev](#)

[Up](#)

[Next](#)

[datatype-pseudo.html](#) [sql.html](#) [functions-logical.html](#)

[8.21. Pseudo-Types](#) [Home](#) [9.1. Logical Operators](#)

Footer

Copyright

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

[Home](#)

[About](#)

[Download](#)

[Documentation](#)

[Community](#)

[Developers](#)

[Support](#)

[Donate](#)

[Your account](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Chapter 10. Type Conversion

Chapter 10. Type Conversion Part II. The SQL Language

[Prev](#)

[Up Home](#)

Chapter 10. Type Conversion

[Table of Contents](#)

- [10.1. Overview](#)
- [10.2. Operators](#)
- [10.3. Functions](#)
- [10.4. Value Storage](#)
- [10.5. `UNION` , `CASE` , and Related Constructs](#)
- [10.6. `SELECT` Output Columns](#)

SQL statements can, intentionally or not, require the mixing of different data types in the same expression. *PostgreSQL* has extensive facilities for evaluating mixed-type expressions.

In many cases a user does not need to understand the details of the type conversion mechanism. However, implicit conversions done by *PostgreSQL* can affect the results of a query. When necessary, these results can be tailored by using *explicit* type conversion.

This chapter introduces the *PostgreSQL* type conversion mechanisms and conventions. Refer to the relevant sections in [Chapter 8](#) and [Chapter 9](#) for more information on specific data types and allowed functions and operators.

[Prev](#)

[Up](#)

[Next](#)

[9.31. Statistics Information Functions](#) [Home](#) [10.1. Overview](#)

PostgreSQL Document Version: 18

Copyright © 1996-2025 The PostgreSQL Global Development Group

Chapter 11. Indexes

Table of Contents

- [11.1. Introduction](#)
- [11.2. Index Types](#)
 - [11.2.1. B-Tree](#)
 - [11.2.2. Hash](#)
 - [11.2.3. GiST](#)
 - [11.2.4. SP-GiST](#)
 - [11.2.5. GIN](#)
 - [11.2.6. BRIN](#)
- [11.3. Multicolumn Indexes](#)
- [11.4. Indexes and *ORDER BY*](#)
- [11.5. Combining Multiple Indexes](#)
- [11.6. Unique Indexes](#)
- [11.7. Indexes on Expressions](#)
- [11.8. Partial Indexes](#)
- [11.9. Index-Only Scans and Covering Indexes](#)
- [11.10. Operator Classes and Operator Families](#)
- [11.11. Indexes and Collations](#)
- [11.12. Examining Index Usage](#)

Indexes are a common way to enhance database performance. An index allows the database server to find and retrieve specific rows much faster than it could do without an index. But indexes also add overhead to the database system as a whole, so they should be used sensibly.

Prev	Up	Next
10.6. SELECT Output Columns	Home	11.1. Introduction

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Part II. The SQL Language

Chapter	Description
Chapter 11. Indexes	Part of PostgreSQL documentation

Indexes are a common way to enhance database performance. An index allows the database server to find and retrieve specific rows much faster than it could do without an index. But indexes also add overhead to the database system as a whole, so they should be used sensibly.

[Previous](#) | [Up](#) | [Next](#)

10.6. SELECT Output Columns | [Home](#) | [11.1. Introduction](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 12. Full Text Search

[Home](#) | [About PostgreSQL](#) | [Privacy Policy](#) | [Code of Conduct](#) | [Contact](#)

News

- May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#)
 - [PostgreSQL 18 Beta 1 Released!](#)
-

Navigation

[Documentation](#) → [PostgreSQL 18](#)

Prev	Up	Next	Home
indexes-examine.html	sql.html	textsearch-intro.html	index.html

Chapter 12. Full Text Search

Table of Contents

- 12.1. Introduction
 - 12.1.1. What Is a Document?
 - 12.1.2. Basic Text Matching
 - 12.1.3. Configurations
- 12.2. Tables and Indexes
 - 12.2.1. Searching a Table
 - 12.2.2. Creating Indexes
- 12.3. Controlling Text Search
 - 12.3.1. Parsing Documents
 - 12.3.2. Parsing Queries
 - 12.3.3. Ranking Search Results
 - 12.3.4. Highlighting Results
- 12.4. Additional Features
 - 12.4.1. Manipulating Documents
 - 12.4.2. Manipulating Queries
 - 12.4.3. Triggers for Automatic Updates
 - 12.4.4. Gathering Document Statistics
- 12.5. Parsers
- 12.6. Dictionaries
 - 12.6.1. Stop Words
 - 12.6.2. Simple Dictionary
 - 12.6.3. Synonym Dictionary
 - 12.6.4. Thesaurus Dictionary
 - 12.6.5. Ispell Dictionary
 - 12.6.6. Snowball Dictionary
- 12.7. Configuration Example
- 12.8. Testing and Debugging Text Search
 - 12.8.1. Configuration Testing

- [12.8.2. Parser Testing](#)
 - [12.8.3. Dictionary Testing](#)
 - [12.9. Preferred Index Types for Text Search](#)
 - [12.10. psql Support](#)
 - [12.11. Limitations](#)
-

Navigation

[Prev](#)

[Up](#)

[Next](#)

[11.12. Examining Index Usage](#) [12.1. Introduction](#) [12.1. Introduction](#)

This documentation is for an unsupported version of PostgreSQL.

You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation → PostgreSQL 18

Supported Versions:

[Current](#) (17) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

[12](#) / [11](#) / [10](#) / [9.6](#) / [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#) / [7.4](#) / [7.3](#) / [7.2](#) / [7.1](#)

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 13. Concurrency Control

Table of Contents

- [13.1. Introduction](#)
- [13.2. Transaction Isolation](#)
 - [13.2.1. Read Committed Isolation Level](#)
 - [13.2.2. Repeatable Read Isolation Level](#)
 - [13.2.3. Serializable Isolation Level](#)
- [13.3. Explicit Locking](#)
 - [13.3.1. Table-Level Locks](#)
 - [13.3.2. Row-Level Locks](#)
 - [13.3.3. Page-Level Locks](#)
 - [13.3.4. Deadlocks](#)
 - [13.3.5. Advisory Locks](#)
- [13.4. Data Consistency Checks at the Application Level](#)
 - [13.4.1. Enforcing Consistency with Serializable Transactions](#)
 - [13.4.2. Enforcing Consistency with Explicit Blocking Locks](#)
- [13.5. Serialization Failure Handling](#)
- [13.6. Caveats](#)
- [13.7. Locking and Indexes](#)

This chapter describes the behavior of the *PostgreSQL* database system when two or more sessions try to access the same data at the same time. The goals in that situation are to allow efficient access for all sessions while maintaining strict data integrity. Every developer of database applications should be familiar with the topics covered in this chapter.

Chapter 13. Concurrency Control

Prev

Up

Home

Next

12.11. Limitations [index.html](#) [mvcc-intro.html](#) 13.1. Introduction

Note: The above content contains links that are placeholders (e.g., `https://example.com/...`) for demonstration purposes.

PostgreSQL: Documentation: 18: Chapter 16. Installation from Binaries

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation → [PostgreSQL 18](#)

Supported Versions:

- [Current \(17 \)](#)
- [16](#)
- [15](#)
- [14](#)

Development Versions:

- [18](#)
 - [devel](#)
-

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 16. Installation from
Binaries

Part III. Server
Administration

[Prev](#)

[Up](#)
[p](#)

[Home](#)

[Ne](#)
[xt](#)

Chapter 16. Installation from Binaries

PostgreSQL is available in the form of binary packages for most common operating systems today. When available, this is the recommended way to install PostgreSQL for users of the system. Building from source (see [Chapter 17](#)) is only recommended for people developing PostgreSQL or extensions.

For an updated list of platforms providing binary packages, please visit the download section on the [PostgreSQL](#) website and follow the instructions for the specific platform.

[Prev](#)

[Up](#) [Next](#)

Part III. Server
Administration

[Home](#) Chapter 17. Installation from Source
[e](#) Code

© 1996-2025 The PostgreSQL Global Development Group | [Privacy Policy](#)
| [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

PostgreSQL: Documentation: 18: Chapter 17. Installation from Source Code

Navigation

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

Latest News

- [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#)
 - [PostgreSQL 18 Beta 1 Released!](#)
-

Breadcrumb Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions

Current: [Current](#) (17) / [16](#) / [15](#) / [14](#) / [13](#) / [Supporting Versions List](#)

Development Versions: [18](#) / [devel](#)

Unsupported versions: [12](#) / [11](#) / [10](#) / [9.6](#) / [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#) / [7.4](#) / [7.3](#) / [7.2](#) / [7.1](#)

Notice

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 17. Installation from Source Code

[Prev](#)

[Up](#)

[Home](#)

[Next](#)

[install-binaries.html](#) [admin.html](#) [index.html](#) [install-requirements.html](#)

Chapter 17. Installation from Source Code

Table of Contents

- [17.1. Requirements](#)
- [17.2. Getting the Source](#)
- [17.3. Building and Installation with Autoconf and Make](#)
 - [17.3.1. Short Version](#)
 - [17.3.2. Installation Procedure](#)

- 17.3.3. [configure](#) Options
- 17.3.4. [configure](#) Environment Variables
- 17.4. Building and Installation with Meson
 - 17.4.1. Short Version
 - 17.4.2. Installation Procedure
 - 17.4.3. [meson setup](#) Options
 - 17.4.4. [meson](#) Build Targets
- 17.5. Post-Installation Setup
 - 17.5.1. Shared Libraries
 - 17.5.2. Environment Variables
- 17.6. Supported Platforms
- 17.7. Platform-Specific Notes
 - 17.7.1. Cygwin
 - 17.7.2. macOS
 - 17.7.3. MinGW
 - 17.7.4. Solaris
 - 17.7.5. Visual Studio

Installation Notice

This chapter describes the installation of *PostgreSQL* using the source code distribution. If you are installing a pre-packaged distribution, such as an RPM or Debian package, ignore this chapter and see [Chapter 16](#) instead.

Navigation

Prev	Up	Next
install-binaries.html	admin.html	install-requirements.html
Chapter	Title	
16	Installation from Binaries	

Chapter 17. Installation from Source Code

This chapter describes the installation of **PostgreSQL** using the source code distribution.

If you are installing a pre-packaged distribution, such as an RPM or Debian package, ignore this chapter and see **Chapter 16** instead.

Table of Contents

1. [17.1. Requirements](#)
 2. [17.2. Getting the Source](#)
 3. [17.3. Building and Installation with Autoconf and Make](#)
 - [17.3.1. Short Version](#)
 - [17.3.2. Installation Procedure](#)
 - [17.3.3. `configure` Options](#)
 - [17.3.4. `configure` Environment Variables](#)
 4. [17.4. Building and Installation with Meson](#)
 - [17.4.1. Short Version](#)
 - [17.4.2. Installation Procedure](#)
 - [17.4.3. `meson setup` Options](#)
 - [17.4.4. `meson` Build Targets](#)
 5. [17.5. Post-Installation Setup](#)
 - [17.5.1. Shared Libraries](#)
 - [17.5.2. Environment Variables](#)
 6. [17.6. Supported Platforms](#)
 7. [17.7. Platform-Specific Notes](#)
-

Source Code Installation

This chapter describes the installation of *PostgreSQL* using the source code distribution. If you are installing a pre-packaged distribution, such as an RPM or Debian package, ignore this chapter and see [Chapter 16](#) instead.

Navigation

Prev

Up

Next

[install-binaries.html](#) [admin.html](#) [install-requirements.html](#)

Footer

Links

- [Privacy Policy](#)
- [Code of Conduct](#)
- [About PostgreSQL](#)
- [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 26. High Availability, Load Balancing, and Replication

[Prev](#) [Up](#) [Home](#) [Next](#)

Chapter 26. High Availability, Load Balancing, and Replication

Table of Contents

- [26.1. Comparison of Different Solutions](#)
- [26.2. Log-Shipping Standby Servers](#)
 - [26.2.1. Planning](#)
 - [26.2.2. Standby Server Operation](#)
 - [26.2.3. Preparing the Primary for Standby Servers](#)
 - [26.2.4. Setting Up a Standby Server](#)
 - [26.2.5. Streaming Replication](#)
 - [26.2.6. Replication Slots](#)
 - [26.2.7. Cascading Replication](#)
 - [26.2.8. Synchronous Replication](#)

- [26.2.9. Continuous Archiving in Standby](#)
 - [26.3. Failover](#)
 - [26.4. Hot Standby](#)
 - [26.4.1. User's Overview](#)
 - [26.4.2. Handling Query Conflicts](#)
 - [26.4.3. Administrator's Overview](#)
 - [26.4.4. Hot Standby Parameter Reference](#)
 - [26.4.5. Caveats](#)
-

Database servers can work together to allow a second server to take over quickly if the primary server fails (high availability), or to allow several computers to serve the same data (load balancing). Ideally, database servers could work together seamlessly. Web servers serving static web pages can be combined quite easily by merely load-balancing web requests to multiple machines. In fact, read-only database servers can be combined relatively easily too. Unfortunately, most database servers have a read/write mix of requests, and read/write servers are much harder to combine. This is because though read-only data needs to be placed on each server only once, a write to any server has to be propagated to all servers so that future read requests to those servers return consistent results.

This synchronization problem is the fundamental difficulty for servers working together. Because there is no single solution that eliminates the impact of the sync problem for all use cases, there are multiple solutions. Each solution addresses this problem in a different way, and minimizes its impact for a specific workload.

Some solutions deal with synchronization by allowing only one server to modify the data. Servers that can modify data are called *read/write*, *master* or *primary* servers. Servers that track changes in the primary are called *standby* or *secondary* servers. A standby server that cannot be connected to until it is promoted to a primary server is called a *warm standby* server, and one that can accept connections and serves read-only queries is called a *hot standby* server.

Some solutions are synchronous, meaning that a data-modifying transaction is not considered committed until all servers have

committed the transaction. This guarantees that a failover will not lose any data and that all load-balanced servers will return consistent results no matter which server is queried. In contrast, asynchronous solutions allow some delay between the time of a commit and its propagation to the other servers, opening the possibility that some transactions might be lost in the switch to a backup server, and that load balanced servers might return slightly stale results. Asynchronous communication is used when synchronous would be too slow.

Solutions can also be categorized by their granularity. Some solutions can deal only with an entire database server, while others allow control at the per-table or per-database level.

Performance must be considered in any choice. There is usually a trade-off between functionality and performance. For example, a fully synchronous solution over a slow network might cut performance by more than half, while an asynchronous one might have a minimal performance impact.

The remainder of this section outlines various failover, replication, and load balancing solutions.

[Prev](#) [Up](#) [Next](#)

Chapter 26. High Availability, Load Balancing, and Replication

Table of Contents

- [26.1. Comparison of Different Solutions](#)
- [26.2. Log-Shipping Standby Servers](#)
 - [26.2.1. Planning](#)
 - [26.2.2. Standby Server Operation](#)
 - [26.2.3. Preparing the Primary for Standby Servers](#)

- [26.2.4. Setting Up a Standby Server](#)
- [26.2.5. Streaming Replication](#)
- [26.2.6. Replication Slots](#)
- [\[26.2.7. Cascading Repli...](#)

PostgreSQL: Documentation: 18: Chapter 27. Monitoring Database Activity

Navigation

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

Latest News

[PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

[Current](#) (17) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

Note

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 27. Monitoring Database Activity

[Prev](#)

[Up](#)

[Home](#)

[Next](#)

[hot-standby.html](#) [admin.html](#) [index.html](#) [monitoring-ps.html](#)

Chapter 27. Monitoring Database Activity

Table of Contents

1. [27.1. Standard Unix Tools](#)
2. [27.2. The Cumulative Statistics System](#)
 - [27.2.1. Statistics Collection Configuration](#)
 - [27.2.2. Viewing Statistics](#)
 - [27.2.3. `pg\_stat\_activity`](#)
 - [27.2.4. `pg\_stat\_replication`](#)
 - [27.2.5. `pg\_stat\_replication\_slots`](#)
 - [27.2.6. `pg\_stat\_wal\_receiver`](#)
 - [27.2.7. `pg\_stat\_recovery\_prefetch`](#)
 - [27.2.8. `pg\_stat\_subscription`](#)
 - [27.2.9. `pg\_stat\_subscription\_stats`](#)

- 27.2.10. `pg_stat_ssl`
- 27.2.11. `pg_stat_gssapi`
- 27.2.12. `pg_stat_archiver`
- 27.2.13. `pg_stat_io`
- 27.2.14. `pg_stat_bgwriter`
- 27.2.15. `pg_stat_checkpoint`
- 27.2.16. `pg_stat_wal`
- 27.2.17. `pg_stat_database`
- 27.2.18. `pg_stat_database_conflicts`
- 27.2.19. `pg_stat_all_tables`
- 27.2.20. `pg_stat_all_indexes`
- 27.2.21. `pg_statio_all_tables`
- 27.2.22. `pg_statio_all_indexes`
- 27.2.23. `pg_statio_all_sequences`
- 27.2.24. `pg_stat_user_functions`
- 27.2.25. `pg_stat_slru`
- 27.2.26. Statistics Functions

3. 27.3. Viewing Locks

4. 27.4. Progress Reporting

- 27.4.1. ANALYZE Progress Reporting
- 27.4.2. CLUSTER Progress Reporting
- 27.4.3. COPY Progress Reporting
- 27.4.4. CREATE INDEX Progress Reporting
- 27.4.5. VACUUM Progress Reporting
- 27.4.6. Base Backup Progress Reporting

5. 27.5. Dynamic Tracing

- 27.5.1. Compiling for Dynamic Tracing
- 27.5.2. Built-in Probes
- 27.5.3. Using Probes
- 27.5.4. Defining New Probes

6. 27.6. Monitoring Disk Usage

- 27.6.1. Determining Disk Usage

Finding Out What the System Is Doing

A database administrator often wonders, “*What is the system doing right now?*” This chapter discusses how to find that out.

Several tools are available for monitoring database activity and analyzing performance. Most of this chapter is devoted to describing PostgreSQL's cumulative statistics system, but one should not neglect regular Unix monitoring programs such as `ps`, `top`, `iostat`, and `vmstat`. Also, once one has identified a poorly-performing query, further investigation might be needed using PostgreSQL's [EXPLAIN](#) command. [Section 14.1](#) discusses `EXPLAIN` and other methods for understanding the behavior of an individual query.

Navigation

[Prev](#) | [Up](#) | [Next](#)

[26.4. Hot Standby](#) | [Home](#) | [27.1. Standard Unix Tools](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

© 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 28. Reliability and the Write-Ahead Log

Navigation

- [Home](#)
- [About](#)
- [Download](#)
- [Documentation](#)
- [Community](#)
- [Developers](#)
- [Support](#)
- [Donate](#)
- [Your account](#)

Latest News

- [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#)
 - [PostgreSQL 18 Beta 1 Released!](#)
-

Documentation Overview

Documentation → PostgreSQL 18

Supported Versions:

- [Current](#) (17), [16](#), [15](#), [14](#), [13](#)*

Development Versions:

- [18](#), [devel](#)

Unsupported Versions:

12 11 10 9.6 9.5 9.4 9.3 9.2 9.1 9.0 8.4 8.3 8.2 8.1 8.0 7.4 7.3 7.2 7.1

Note

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 28. Reliability and the Write-Ahead Log

< Prev	Up	Home	Next >
diskusage.html	admin.html	Part III. Server Administration	wal-reliability.html

Chapter Overview

Chapter 28. Reliability and the Write-Ahead Log

[Table of Contents](#)

- [28.1. Reliability](#)
- [28.2. Data Checksums](#)
 - [28.2.1. Off-line Enabling of Checksums](#)
- [28.3. Write-Ahead Logging \(WAL\)](#)
- [28.4. Asynchronous Commit](#)
- [28.5. WAL Configuration](#)
- [28.6. WAL Internals](#)

This chapter explains how to control the reliability of PostgreSQL, including details about the Write-Ahead Log.

Navigation

[< Prev](#) [Up](#) [Next >](#)

[diskusage.html](#) [admin.html](#) [wal-reliability.html](#)

[27.6. Monitoring Disk Usage](#) [Home](#) [28.1. Reliability](#)

[Previous](#)

[Next](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

© 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 29. Logical Replication

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Chapter 29. Logical Replication

[Prev Up](#) [Part III. Server Administration](#) [Home](#) [Next](#)

Chapter 29. Logical Replication

Table of Contents

- [29.1. Publication](#)
 - [29.1.1.. Replica Identity](#)
- [29.2. Subscription](#)
 - [29.2.1.. Replication Slot Management](#)
 - [29.2.2.. Examples: Set Up Logical Replication](#)

- 29.2.3.. Examples: Deferred Replication Slot Creation
- 29.3.. Logical Replication Failover
- 29.4.. Row Filters
 - 29.4.1.. Row Filter Rules
 - 29.4.2.. Expression Restrictions
 - 29.4.3.. UPDATE Transformations
 - 29.4.4.. Partitioned Tables
 - 29.4.5.. Initial Data Synchronization
 - 29.4.6.. Combining Multiple Row Filters
 - 29.4.7.. Examples
- 29.5.. Column Lists
 - 29.5.1.. Examples
- 29.6.. Generated Column Replication
- 29.7.. Conflicts
- 29.8.. Restrictions
- 29.9.. Architecture
 - 29.9.1.. Initial Snapshot
- 29.10.. Monitoring
- 29.11.. Security
- 29.12.. Configuration Settings
 - 29.12.1.. Publishers
 - 29.12.2.. Subscribers
- 29.13.. Upgrade
 - 29.13.1.. Prepare for publisher upgrades
 - 29.13.2.. Prepare for subscriber upgrades
 - 29.13.3.. Upgrading logical replication clusters
- 29.14.. Quick Setup

Logical replication is a method of replicating data objects and their changes, based upon their **replication identity** (usually a primary key). We use the term *logical* in contrast to **physical** replication, which uses exact block addresses and byte-by-byte replication. PostgreSQL supports both mechanisms concurrently, see [Chapter 26](#). Logical replication allows fine-grained control over both data replication and security.

Logical replication uses a *publish* and *subscribe* model with one or more *subscribers* subscribing to one or more *publications* on a *publisher* node.

Subscribers pull data from the publications they subscribe to and may subsequently re-publish data to allow cascading replication or more complex configurations.

When logical replication of a table typically starts, PostgreSQL takes a snapshot of the table's data on the publisher database and copies it to the subscriber. Once complete, changes on the publisher since the initial copy are sent continually to the subscriber. The subscriber applies the data in the same order as the publisher so that transactional consistency is guaranteed for publications within a single subscription. This method of data replication is sometimes referred to as transactional replication.

The typical use-cases for logical replication are:

- Sending incremental changes in a single database or a subset of a database to subscribers as they occur.
- Firing triggers for individual changes as they arrive on the subscriber.
- Consolidating multiple databases into a single one (for example for analytical purposes).
- Replicating between different major versions of PostgreSQL.
- Replicating between PostgreSQL instances on different platforms (for example Linux to Windows).
- Giving access to replicated data to different groups of users.
- Sharing a subset of the database between multiple databases.

The subscriber database behaves in the same way as any other PostgreSQL instance and can be used as a publisher for other databases by defining its own publications. When the subscriber is treated as read-only by application, there will be no conflicts from a single subscription. On the other hand, if there are other writes done either by an application or by other subscribers to the same set of tables, conflicts can arise.

[Prev](#) | [Up](#) | [Next](#)

28.6. WAL Internals | [Home](#) | 29.1. Publication

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 30. Just-in- Time Compilation (JIT)

Navigation

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

Latest News

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

[Current](#) (17) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

[12 / 11](#)

Notice

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 30. Just-in-Time Compilation (JIT)

Prev	Up	Part III. Server Administration	Home	Next
logical-replication-quick-setup.html	admin.html	Part III. Server Administration	index.html	jit-reason.html

Chapter 30. Just-in-Time Compilation (JIT)

Table of Contents

- [30.1. What Is JIT compilation?](#)
 - [30.1.1. JIT Accelerated Operations](#)
 - [30.1.2. Inlining](#)
 - [30.1.3. Optimization](#)
- [30.2. When to JIT?](#)
- [30.3. Configuration](#)
- [30.4. Extensibility](#)
 - [30.4.1. Inlining Support for Extensions](#)
 - [30.4.2. Pluggable JIT Providers](#)

This chapter explains what just-in-time compilation is, and how it can be configured in *PostgreSQL*.

Navigation

[Prev](#) | [Up](#) | [Next](#)

| [29.14. Quick Setup](#) | [Home](#) | [30.1. What Is JIT compilation?](#) |

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996–2025 The PostgreSQL Global
Development Group

[Home](#)
[About](#)
[Download](#)
[Documentation](#)
[Community](#)
[Developers](#)
[Support](#)
[Donate](#)
[Your account](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Chapter 18. Server Setup and Operation

Table of Contents

- [18.1. The PostgreSQL User Account](#)
- [18.2. Creating a Database Cluster](#)
 - [18.2.1. Use of Secondary File Systems](#)
 - [18.2.2. File Systems](#)
- [18.3. Starting the Database Server](#)
 - [18.3.1. Server Start-up Failures](#)
 - [18.3.2. Client Connection Problems](#)
- [18.4. Managing Kernel Resources](#)
 - [18.4.1. Shared Memory and Semaphores](#)
 - [18.4.2. systemd RemoveIPC](#)
 - [18.4.3. Resource Limits](#)
 - [18.4.4. Linux Memory Overcommit](#)

- [18.4.5. Linux Huge Pages](#)
- [18.5. Shutting Down the Server](#)
- [18.6. Upgrading a PostgreSQL Cluster](#)
 - [18.6.1. Upgrading Data via pg_dumpall](#)
 - [18.6.2. Upgrading Data via pg_upgrade](#)
 - [18.6.3. Upgrading Data via Replication](#)
- [18.7. Preventing Server Spoofing](#)
- [18.8. Encryption Options](#)
- [18.9. Secure TCP/IP Connections with SSL](#)
 - [18.9.1. Basic Setup](#)
 - [18.9.2. OpenSSL Configuration](#)
 - [18.9.3. Using Client Certificates](#)
 - [18.9.4. SSL Server File Usage](#)
 - [18.9.5. Creating Certificates](#)
- [18.10. Secure TCP/IP Connections with GSSAPI Encryption](#)
 - [18.10.1. Basic Setup](#)
- [18.11. Secure TCP/IP Connections with SSH Tunnels](#)
- [18.12. Registering Event Log on Windows](#)

This chapter discusses how to set up and run the database server, and its interactions with the operating system.

The directions in this chapter assume that you are working with plain *PostgreSQL* without any additional infrastructure, for example a copy that you built from source according to the directions in the preceding chapters. If you are working with a pre-packaged or vendor-supplied version of *PostgreSQL*, it is likely that the packager has made special provisions for installing and starting the database server according to your system's conventions. Consult the package-level documentation for details.

[Prev](#) | [Up](#) | [Next](#)

17.7. Platform-Specific Notes

This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of

the other supported versions listed above instead.

Chapter 18. Server Setup and Operation		Part III. Server Administration
Prev	Up Home	Next

Chapter 18. Server Setup and Operation

Table of Contents

- [18.1. The PostgreSQL User Account](#)
- [18.2. Creating a Database Cluster](#)
 - [18.2.1. Use of Secondary File Systems](#)
 - [18.2.2. File Systems](#)
- [18.3. Starting the Database Server](#)
 - [18.3.1. Server Start-up Failures](#)
 - [18.3.2. Client Connection Problems](#)
- [18.4. Managing Kernel Resources](#)
 - [18.4.1. Shared Memory and Semaphores](#)
 - [18.4.2. systemd RemoveIPC](#)
 - [18.4.3. Resource Limits](#)
 - [18.4.4. Linux Memory Overcommit](#)
 - [18.4.5. Linux Huge Pages](#)
- [18.5. Shutting Down the Server](#)
- [18.6. Upgrading a PostgreSQL Cluster](#)
 - [18.6.1. Upgrading Data via pg_dumpall](#)
 - [18.6.2. Upgrading Data via pg_upgrade](#)
 - [18.6.3. Upgrading Data via Replication](#)
- [18.7. Preventing Server Spoofing](#)
- [18.8. Encryption Options](#)
- [18.9. Secure TCP/IP Connections with SSL](#)
 - [18.9.1. Basic Setup](#)
 - [18.9.2. OpenSSL Configuration](#)
 - [18.9.3. Using Client Certificates](#)
 - [18.9.4. SSL Server File Usage](#)
 - [18.9.5. Creating Certificates](#)
- [18.10. Secure TCP/IP Connections with GSSAPI Encryption](#)

- [18.10.1. Basic Setup](#)
- [18.11. Secure TCP/IP Connections with SSH Tunnels](#)
- [18.12. Registering Event Log on Windows](#)

This chapter discusses how to set up and run the database server, and its interactions with the operating system.

The directions in this chapter assume that you are working with plain *PostgreSQL* without any additional infrastructure, for example a copy that you built from source according to the directions in the preceding chapters. If you are working with a pre-packaged or vendor-supplied version of *PostgreSQL*, it is likely that the packager has made special provisions for installing and starting the database server according to your system's conventions. Consult the package-level documentation for details.

[Prev](#) | [Up](#) | [Next](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 19. Server Configuration

Navigation Menu

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

News

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Breadcrumb Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions

Current:

[Current](#), (17), 16, 15, 14, 13

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

12, 11, 10, 9.6, 9.5, 9.4, 9.3, 9.2, 9.1, 9.0, 8.4, 8.3, 8.2, 8.1, 8.0, 7.4, 7.3, 7.2, 7.1

Note

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 19. Server Configuration

Navigation

[Prev](#)

[Up](#)

[Next](#)

[Event Log Registration on Windows Admin Setting Parameters](#)

Chapter Title

Chapter 19. Server Configuration

Table of Contents

1. 19.1. Setting Parameters
 - 19.1.1. Parameter Names and Values
 - 19.1.2. Parameter Interaction via the Configuration File
 - 19.1.3. Parameter Interaction via SQL
 - 19.1.4. Parameter Interaction via the Shell
 - 19.1.5. Managing Configuration File Contents
2. 19.2. File Locations
3. 19.3. Connections and Authentication
 - 19.3.1. Connection Settings
 - 19.3.2. TCP Settings
 - 19.3.3. Authentication
 - 19.3.4. SSL
4. 19.4. Resource Consumption
 - 19.4.1. Memory
 - 19.4.2. Disk
 - 19.4.3. Kernel Resource Usage
 - 19.4.4. Background Writer
 - 19.4.5. I/O
 - 19.4.6. Worker Processes
5. 19.5. Write Ahead Log
 - 19.5.1. Settings
 - 19.5.2. Checkpoints
 - 19.5.3. Archiving
 - 19.5.4. Recovery
 - 19.5.5. Archive Recovery
 - 19.5.6. Recovery Target
 - 19.5.7. WAL Summarization
6. 19.6. Replication
 - 19.6.1. Sending Servers
 - 19.6.2. Primary Server
 - 19.6.3. Standby Servers
 - 19.6.4. Subscribers
7. 19.7. Query Planning

- 19.7.1. Planner Method Configuration
 - 19.7.2. Planner Cost Constants
 - 19.7.3. Genetic Query Optimizer
 - 19.7.4. Other Planner Options
 - 8. 19.8. Error Reporting and Logging
 - 19.8.1. Where to Log
 - 19.8.2. When to Log
 - 19.8.3. What to Log
 - 19.8.4. Using CSV-Format Log Output
 - 19.8.5. Using JSON-Format Log Output
 - 19.8.6. Process Title
 - 9. 19.9. Run-time Statistics
 - 19.9.1. Cumulative Query and Index Statistics
 - 19.9.2. Statistics Monitoring
 - 10. 19.10. Vacuuming
 - 19.10.1. Automatic Vacuuming
 - 19.10.2. Cost-based Vacuum Delay
 - 19.10.3. Default Behavior
 - 19.10.4. Freezing
 - 11. 19.11. Client Connection Defaults
 - 19.11.1. Statement Behavior
 - 19.11.2. Locale and Formatting
 - 19.11.3. Shared Library Preloading
 - 19.11.4. Other Defaults
 - 12. 19.12. Lock Management
 - 13. 19.13. Version and Platform Compatibility
 - 19.13.1. Previous PostgreSQL Versions
 - 19.13.2. Platform and Client Compatibility
 - 14. 19.14. Error Handling
 - 15. 19.15. Preset Options
 - 16. 19.16. Customized Options
 - 17. 19.17. Developer Options
 - 18. 19.18. Short Options
-

Paragraph

There are many configuration parameters that affect the behavior of the database system. In the first section of this chapter, we describe how to interact with configuration parameters. The subsequent sections discuss each parameter in detail.

Navigation

Prev

Up

Next

[Registering Event Log on Windows Home Setting Parameters](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 20. Client Authentication

[PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation → [PostgreSQL 18](#)

Supported Versions:

[Current \(17\)](#) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

[12](#) / [11](#) / [10](#) / [9.6](#) / [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#) / [7.4](#) / [7.3](#) / [7.2](#) / [7.1](#)

This documentation is for an unsupported version of PostgreSQL.

You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 20. Client Authentication

Table of Contents

- [20.1. The `pg\_hba.conf` File](#)
- [20.2. User Name Maps](#)
- [20.3. Authentication Methods](#)
- [20.4. Trust Authentication](#)
- [20.5. Password Authentication](#)
- [20.6. GSSAPI Authentication](#)
- [20.7. SSPI Authentication](#)
- [20.8. Ident Authentication](#)
- [20.9. Peer Authentication](#)
- [20.10. LDAP Authentication](#)
- [20.11. RADIUS Authentication](#)
- [20.12. Certificate Authentication](#)
- [20.13. PAM Authentication](#)
- [20.14. BSD Authentication](#)
- [20.15. OAuth Authorization/Authentication](#)
- [20.16. Authentication Problems](#)

When a client application connects to the database server...

It specifies which PostgreSQL database user name it wants to connect as, much the same way one logs into a Unix computer as a particular user. Within the SQL environment, the active database user name

determines access privileges to database objects — see [Chapter 21](#) for more information. Therefore, it is essential to restrict which database users can connect.

Note

As explained in [Chapter 21](#), PostgreSQL actually manages privileges in terms of *roles*. In this chapter, we use *database user* to mean *role with the `LOGIN` privilege*.

Authentication is the process by which the database server establishes the identity of the client, and by extension determines whether the client application (or the user who runs the client application) is permitted to connect with the requested database user name.

PostgreSQL offers a number of different client authentication methods. The method used to authenticate a particular client connection can be selected based on (client) host address, database, and user.

PostgreSQL database user names are logically separate from user names of the operating system in which the server runs. If all the users of a particular server also have accounts on the server's machine, it makes sense to assign database user names that match their operating system user names. However, a server that accepts remote connections might have many database users who have no local OS account, and in such cases, there need be no connection between database user names and OS user names.

Prev		Up	Next	
runtime-config-short.html		admin.html	auth-pg-hba-conf.html	
Short Options	Home	20.1. The	<code>pg_hba.conf</code>	File
19.18	index.html	20.1. The	<code>pg_hba.conf</code>	File

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 21. Database Roles

Navigation

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

News

[PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Breadcrumbs

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

[Current](#) (17), [16](#), [15](#), [14](#), [13](#)

Development Versions:

[18](#), [devel](#)

Unsupported versions:

[12](#), [11](#), [10](#), [9.6](#), [9.5](#), [9.4](#), [9.3](#), [9.2](#), [9.1](#), [9.0](#), [8.4](#), [8.3](#), [8.2](#), [8.1](#), [8.0](#), [7.4](#), [7.3](#), [7.2](#), [7.1](#)

Warning

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 21. Database Roles

Prev	Up	Part III. Server Administration	Home	Next
client-authentication-problems.html	admin.html	<i>Part III</i>	index.html	database-roles.html

Chapter 21. Database Roles

Table of Contents

- [21.1. Database Roles](#)
 - [21.2. Role Attributes](#)
 - [21.3. Role Membership](#)
 - [21.4. Dropping Roles](#)
 - [21.5. Predefined Roles](#)
 - [21.6. Function Security](#)
-

PostgreSQL manages database access permissions using the concept of *roles*. A role can be thought of as either a database *user*, or a *group* of database users, depending on how the role is set up. Roles can own database objects (for example, tables and functions) and can assign privileges on those objects to other roles to control who has access to which objects. Furthermore, it is possible to grant *membership* in a role

to another role, thus allowing the member role to use privileges assigned to another role.

The concept of roles subsumes the concepts of *users* and *groups*. In PostgreSQL versions before 8.1, users and groups were distinct kinds of entities, but now there are only roles. Any role can act as a user, a group, or both.

This chapter describes how to create and manage roles. More information about the effects of role privileges on various database objects can be found in [Section 5.8](#).

Navigation

[Prev](#) | [Up](#) | [Next](#)

[20.16. Authentication Problems](#) | [Home](#) | [21.1. Database Roles](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL:

Documentation: 18:

Chapter 22. Managing Databases

- [Home](#)
- [About](#)
- [Download](#)
- [Documentation](#)
- [Community](#)
- [Developers](#)
- [Support](#)
- [Donate](#)
- [Your account](#)

```
<div>
  <span>
    <i></i>
  </span>
</div>
</div>
```

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions: - [Current]

(<https://www.postgresql.org/docs/current/managing-databases.html>)

(17) - [16]([https://www.postgresql.org/docs/16/managing-](https://www.postgresql.org/docs/16/managing-databases.html)

[databases.html](https://www.postgresql.org/docs/16/managing-databases.html)) - [15]([https://www.postgresql.org/docs/15/managing-](https://www.postgresql.org/docs/15/managing-databases.html)

[databases.html](https://www.postgresql.org/docs/15/managing-databases.html)) - [14]([111](https://www.postgresql.org/docs/14/managing-</p></div><div data-bbox=)

databases.html) - [13](https://www.postgresql.org/docs/13/managing-databases.html)

Development Versions: - [18]

(https://www.postgresql.org/docs/18/managing-databases.html) - [devel]

(https://www.postgresql.org/docs/devel/managing-databases.html)

Unsupported versions: - [12]

(https://www.postgresql.org/docs/12/managing-databases.html) - [11]

(https://www.postgresql.org/docs/11/managing-databases.html) - [10]

(https://www.postgresql.org/docs/10/managing-databases.html) - [9.6]

(https://www.postgresql.org/docs/9.6/managing-databases.html) - [9.5]

(https://www.postgresql.org/docs/9.5/managing-databases.html) - [9.4]

(https://www.postgresql.org/docs/9.4/managing-databases.html) - [9.3]

(https://www.postgresql.org/docs/9.3/managing-databases.html) - [9.2]

(https://www.postgresql.org/docs/9.2/managing-databases.html) - [9.1]

(https://www.postgresql.org/docs/9.1/managing-databases.html) - [9.0]

(https://www.postgresql.org/docs/9.0/managing-databases.html) - [8.4]

(https://www.postgresql.org/docs/8.4/managing-databases.html) - [8.3]

(https://www.postgresql.org/docs/8.3/managing-databases.html) - [8.2]

(https://www.postgresql.org/docs/8.2/managing-databases.html) - [8.1]

(https://www.postgresql.org/docs/8.1/managing-databases.html) - [8.0]

(https://www.postgresql.org/docs/8.0/managing-databases.html) - [7.4]

(https://www.postgresql.org/docs/7.4/managing-databases.html) - [7.3]

(https://www.postgresql.org/docs/7.3/managing-databases.html) - [7.2]

(https://www.postgresql.org/docs/7.2/managing-databases.html) - [7.1]

(https://www.postgresql.org/docs/7.1/managing-databases.html)

This documentation is for an unsupported version of PostgreSQL.

You may want to view the same page for the [current]

(https://www.postgresql.org/docs/current/managing-databases.html)

version, or one of the other supported versions listed above instead.

[Prev Up](#) [Part III. Server Administration](#) [Home](#) [Next](#)

Chapter 22. Managing Databases

Table of Contents

- [22.1. Overview](#)
- [22.2. Creating a Database](#)
- [22.3. Template Databases](#)
- [22.4. Database Configuration](#)
- [22.5. Destroying a Database](#)
- [22.6. Tablespaces](#)

Every instance of a running PostgreSQL server manages one or more databases. Databases are therefore the topmost hierarchical level for organizing SQL objects (“database objects”). This chapter describes the properties of databases, and how to create, manage, and destroy them.

[Prev](#) [Up](#) [Next](#)

[21.6. Function Security](#) [Home](#) [22.1. Overview](#)

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

Chapter 23. Localization

Table of Contents

23.1. Locale Support

23.1.1. Overview

23.1.2. Behavior

23.1.3. Selecting Locales

23.1.4. Locale Providers

23.1.5. ICU Locales

23.1.6. Problems

23.2. Collation Support

23.2.1. Concepts

23.2.2. Managing Collations

23.2.3. ICU Custom Collations

23.3. Character Set Support

23.3.1. Supported Character Sets

23.3.2. Setting the Character Set

23.3.3. Automatic Character Set Conversion Between Server and Client

23.3.4. Available Character Set Conversions

23.3.5. Further Reading

This chapter describes the available localization features from the point of view of the administrator. *PostgreSQL* supports two localization facilities:

- Using the locale features of the operating system to provide locale-specific collation order, number formatting, translated messages, and other aspects. This is covered in [Section 23.1](#) and [Section 23.2](#).
- Providing a number of different character sets to support storing text in all kinds of languages, and providing character set translation between client and server. This is covered in [Section 23.3](#).



This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

[Prev](#)

[Up](#)

[Next](#)

[manage-ag-tablespaces.html](#) [admin.html](#) [locale.html](#)

[22.6. Tablespaces](#) [23.1. Locale Support](#)

Chapter 23. Localization

Table of Contents

23.1. Locale Support

23.1.1. Overview

23.1.2. Behavior

23.1.3. Selecting Locales

23.1.4. Locale Providers

23.1.5. ICU Locales

23.1.6. Problems

23.2. Collation Support

23.2.1. Concepts

23.2.2. Managing Collations

23.2.3. ICU Custom Collations

23.3. Character Set Support

23.3.1. Supported Character Sets

23.3.2. Setting the Character Set

23.3.3. Automatic Character Set Conversion Between Server and Client

23.3.4. Available Character Set Conversions

23.3.5. Further Reading

PostgreSQL supports two localization facilities:

- Using the locale features of the operating system to provide locale-specific collation order, number formatting, translated messages, and other aspects. This is covered in [Section 23.1](#) and [Section 23.2](#).

- Providing a number of different character sets to support storing text in all kinds of languages, and providing character set translation between client and server. This is covered in [Section 23.3](#).
-

Additional Information

This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Navigation:

- [Prev](#)
- [Up](#)
- [Next](#)

PostgreSQL: Documentation: 18: Chapter 24. Routine Database Maintenance Tasks

Navigation

[Prev](#)

[Up](#)

[Home](#)

[Next](#)

[multibyte.html](#) [admin.html](#) [index.html](#) [routine-vacuuming.html](#)

Chapter 24. Routine Database Maintenance Tasks

Table of Contents

- [24.1. Routine Vacuuming](#)
 - [24.1.1. Vacuuming Basics](#)
 - [24.1.2. Recovering Disk Space](#)
 - [24.1.3. Updating Planner Statistics](#)
 - [24.1.4. Updating the Visibility Map](#)
 - [24.1.5. Preventing Transaction ID Wraparound Failures](#)
 - [24.1.6. The Autovacuum Daemon](#)
- [24.2. Routine Reindexing](#)
- [24.3. Log File Maintenance](#)

Important Notice

PostgreSQL, like any database software, requires that certain tasks be performed regularly to achieve optimum performance. The tasks discussed here are *required*, but they are repetitive in nature and can easily be automated using standard tools such as *cron* scripts or Windows' *Task Scheduler*. It is the database administrator's responsibility to set up appropriate scripts, and to check that they execute successfully.

One obvious maintenance task is the creation of backup copies of the data on a regular schedule. Without a recent backup, you have no chance of recovery after a catastrophe (disk failure, fire, mistakenly dropping a critical table, etc.). The backup and recovery mechanisms available in PostgreSQL are discussed at length in [Chapter 25](#).

The other main category of maintenance task is periodic “*vacuuming*” of the database. This activity is discussed in [Section 24.1](#). Closely related to this is updating the statistics that will be used by the query planner, as discussed in [Section 24.1.3](#).

Another task that might need periodic attention is log file management. This is discussed in [Section 24.3](#).

[check_postgres](#) is available for monitoring database health and reporting unusual conditions. `check_postgres` integrates with Nagios and MRTG, but can be run standalone too.

PostgreSQL is low-maintenance compared to some other database management systems. Nonetheless, appropriate attention to these tasks will go far towards ensuring a pleasant and productive experience with the system.

Navigation (continued)

Prev

Up

Next

23.3. Character Set Support [index.html](#) 24.1. Routine Vacuuming

Prev

Up

Next

[multibyte.html](#) [admin.html](#) [routine-vacuuming.html](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 25. Backup and Restore

Navigation Menu

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

News

- [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)
-

Navigation Breadcrumbs

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)
- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions:

- [18](#)
- [devel](#)

Unsupported Versions:

- [12](#)
 - [11](#)
 - [10](#)
 - [9.6](#)
 - [9.5](#)
 - [9.4](#)
 - [9.3](#)
 - [9.2](#)
 - [9.1](#)
 - [9.0](#)
 - [8.4](#)
 - [8.3](#)
 - [8.2](#)
 - [8.1](#)
 - [8.0](#)
 - [7.4](#)
 - [7.3](#)
 - [7.2](#)
 - [7.1](#)
-

Disclaimer

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 25. Backup and Restore

Chapter 25. Backup
and Restore

Prev

Up

Home

Next

[logfile-
maintenance.ht
ml](#)

[admin.
html](#)

[index.ht
ml](#)

[backup-
dump.html](#)

Chapter 25. Backup and Restore Content

Table of Contents

- [25.1. SQL Dump](#)
 - [25.1.1. Restoring the Dump](#)
 - [25.1.2. Using pg_dumpall](#)
 - [25.1.3. Handling Large Databases](#)
- [25.2. File System Level Backup](#)
- [\[25.3. Continuous Archiving and Point-in-Time Recovery \(PITR\)\]](#)
[continuous-archiving.html](#)
 - [25.3.1. Setting Up WAL Archiving](#)
 - [25.3.2. Making a Base Backup](#)
 - [25.3.3. Making an Incremental Backup](#)
 - [25.3.4. Making a Base Backup Using the Low Level API](#)
 - [25.3.5. Recovering Using a Continuous Archive Backup](#)

- [25.3.6. Timelines](#)
 - [25.3.7. Tips and Examples](#)
 - [25.3.8. Caveats](#)
-

Important Notice

As with everything that contains valuable data, *PostgreSQL* databases should be backed up regularly. While the procedure is essentially simple, it is important to have a clear understanding of the underlying techniques and assumptions.

Approaches to Backing Up PostgreSQL Data:

- SQL dump
- File system level backup
- Continuous archiving

Each has its own strengths and weaknesses; each is discussed in turn in the following sections.

Navigation

[Prev](#)

[Up](#)

[Next](#)

[24.3. Log File Maintenance index.html](#) [25.1. SQL Dump](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

© 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Legal Notice

[Home](#)
[About](#)
[Download](#)
[Documentation](#)
[Community](#)
[Developers](#)
[Support](#)
[Donate](#)
[Your account](#)

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation for PostgreSQL 18

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)
- [16](#)
- [15](#)
- [14](#)

- [13](#)

Development Versions:

- [18](#)
- [devel](#)

Unsupported versions:

- [12](#)
- [11](#)
- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)
- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)
- [8.1](#)
- [8.0](#)
- [7.4](#)

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of
the other supported versions listed above instead.

Legal Notice

PostgreSQL is Copyright © 1996–2025 by the PostgreSQL Global
Development Group.

Postgres95 is Copyright © 1994–5 by the Regents of the University of
California.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN “AS-IS” BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group



Chapter 35. The Information Schema

[Prev](#) [Up](#) [Part IV. Client Interfaces](#) [Home](#) [Next](#)

Chapter 35. The Information Schema

Table of Contents

- [35.1. The Schema](#)
- [35.2. Data Types](#)
- [35.3. information_schema_catalog_name](#)
- [35.4. administrable_role_authorizations](#)
- [35.5. applicable_roles](#)
- [35.6. attributes](#)
- [35.7. character_sets](#)
- [35.8. check_constraint_routine_usage](#)
- [35.9. check_constraints](#)
- [35.10. collations](#)
- [35.11. collation_character_set_applicability](#)
- [35.12. column_column_usage](#)
- [35.13. column_domain_usage](#)
- [35.14. column_options](#)
- [35.15. column_privileges](#)
- [35.16. column_udt_usage](#)
- [35.17. columns](#)
- [35.18. constraint_column_usage](#)

- 35.19. `constraint_table_usage`
- 35.20. `data_type_privileges`
- 35.21. `domain_constraints`
- 35.22. `domain_udt_usage`
- 35.23. `domains`
- 35.24. `element_types`
- 35.25. `enabled_roles`
- 35.26. `foreign_data_wrapper_options`
- 35.27. `foreign_data_wrappers`
- 35.28. `foreign_server_options`
- 35.29. `foreign_servers`
- 35.30. `foreign_table_options`
- 35.31. `foreign_tables`
- 35.32. `key_column_usage`
- 35.33. `parameters`
- 35.34. `referential_constraints`
- 35.35. `role_column_grants`
- 35.36. `role_routine_grants`
- 35.37. `role_table_grants`
- 35.38. `role_udt_grants`
- 35.39. `role_usage_grants`
- 35.40. `routine_column_usage`
- 35.41. `routine_privileges`
- 35.42. `routine_routine_usage`
- 35.43. `routine_sequence_usage`
- 35.44. `routine_table_usage`
- 35.45. `routines`
- 35.46. `schemata`
- 35.47. `sequences`
- 35.48. `sql_features`
- 35.49. `sql_implementation_info`
- 35.50. `sql_parts`
- 35.51. `sql_sizing`
- 35.52. `table_constraints`
- 35.53. `table_privileges`
- 35.54. `tables`
- 35.55. `transforms`

- 35.56. `triggered_update_columns`
- 35.57. `triggers`
- 35.58. `udt_privileges`
- 35.59. `usage_privileges`
- 35.60. `user_defined_types`
- 35.61. `user_mapping_options`
- 35.62. `user_mappings`
- 35.63. `view_column_usage`
- 35.64. `view_routine_usage`
- 35.65. `view_table_usage`
- 35.66. `views`

The information schema consists of a set of views that contain information about the objects defined in the current database. The information schema is defined in the SQL standard and can therefore be expected to be portable and remain stable — unlike the system catalogs, which are specific to *PostgreSQL* and are modeled after implementation concerns. The information schema views do not, however, contain information about *PostgreSQL*-specific features; to inquire about those you need to query the system catalogs or other *PostgreSQL*-specific views.

Note

When querying the database for constraint information, it is possible for a standard-compliant query that expects to return one row to return several. This is because the SQL standard requires constraint names to be unique within a schema, but *PostgreSQL* does not enforce this restriction. *PostgreSQL* automatically-generated constraint names avoid duplicates in the same schema, but users can specify such duplicate names.

This problem can appear when querying information schema views such as `check_constraint_routine_usage`, `check_constraints`, `domain_constraints`, and `referential_constraints`. Some other views have similar issues but contain the table name to help distinguish duplicate rows, e.g.,

```
constraint_column_usage , constraint_table_usage ,  
table_constraints .
```

Navigation

Prev

Up

Next

[34.17. Internals](#) [Home](#) [35.1. The Schema](#)

PostgreSQL: Documentation: 18: Chapter 36. Extending SQL

Navigation

[Home](#) | [About](#) | [Download](#) | [Documentation](#) | [Community](#) | [Developers](#) | [Support](#) | [Donate](#) | [Your account](#)

News

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation for PostgreSQL 18

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

[Current](#) (17) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

Unsupported versions:

[12](#) / [11](#) / [10](#) / [9.6](#) / [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#) / [7.4](#) / [7.3](#) / [7.2](#) / [7.1](#)

Note:

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 36. Extending SQL

Prev	Up	Part V. Server Programming	Home	Next
server-programming.html	server-programming.html	Part V. Server Programming	index.html	extend-how.html

Chapter 36. Extending SQL

Table of Contents

- [36.1. How Extensibility Works](#)
- [36.2. The PostgreSQL Type System](#)
 - 36.2.1. Base Types
 - 36.2.2. Container Types
 - 36.2.3. Domains
 - 36.2.4. Pseudo-Types
 - 36.2.5. Polymorphic Types
- [36.3. User-Defined Functions](#)

- [36.4. User-Defined Procedures](#)
- [36.5. Query Language \(SQL\) Functions](#)
 - 36.5.1. Arguments for SQL Functions
 - 36.5.2. SQL Functions on Base Types
 - 36.5.3. SQL Functions on Composite Types
 - 36.5.4. SQL Functions with Output Parameters
 - 36.5.5. SQL Procedures with Output Parameters
 - 36.5.6. SQL Functions with Variable Numbers of Arguments
 - 36.5.7. SQL Functions with Default Values for Arguments
 - 36.5.8. SQL Functions as Table Sources
 - 36.5.9. SQL Functions Returning Sets
 - 36.5.10. SQL Functions Returning `TABLE`
 - 36.5.11. Polymorphic SQL Functions
 - 36.5.12. SQL Functions with Collations
- [36.6. Function Overloading](#)
- [36.7. Function Volatility Categories](#)
- [36.8. Procedural Language Functions](#)
- [36.9. Internal Functions](#)
- [36.10. C-Language Functions](#)
 - 36.10.1. Dynamic Loading
 - 36.10.2. Base Types in C-Language Functions
 - 36.10.3. Version 1 Calling Conventions
 - 36.10.4. Writing Code
 - 36.10.5. Compiling and Linking Dynamically-Loaded Functions
 - 36.10.6. Server API and ABI Stability Guidance
 - 36.10.7. Composite-Type Arguments
 - 36.10.8. Returning Rows (Composite Types)
 - 36.10.9. Returning Sets
 - 36.10.10. Polymorphic Arguments and Return Types
 - 36.10.11. Shared Memory
 - 36.10.12. LWLocks

- 36.10.13. Custom Wait Events
- 36.10.14. Injection Points
- 36.10.15. Custom Cumulative Statistics
- 36.10.16. Using C++ for Extensibility
- [36.11. Function Optimization Information](#)
- [36.12. User-Defined Aggregates](#)
 - 36.12.1. Moving-Aggregate Mode
 - 36.12.2. Polymorphic and Variadic Aggregates
 - 36.12.3. Ordered-Set Aggregates
 - 36.12.4. Partial Aggregation
 - 36.12.5. Support Functions for Aggregates
- [36.13. User-Defined Types](#)
 - 36.13.1. TOAST Considerations
- [36.14. User-Defined Operators](#)
- [36.15. Operator Optimization Information](#)
 - 36.15.1. COMMUTATOR
 - 36.15.2. NEGATOR
 - 36.15.3. RESTRICT
 - 36.15.4. JOIN
 - 36.15.5. HASHES
 - 36.15.6. MERGES
- [36.16. Interfacing Extensions to Indexes](#)
 - 36.16.1. Index Methods and Operator Classes
 - 36.16.2. Index Method Strategies
 - 36.16.3. Index Method Support Routines
 - 36.16.4. An Example
 - 36.16.5. Operator Classes and Operator Families
 - 36.16.6. System Dependencies on Operator Classes
 - 36.16.7. Ordering Operators
 - 36.16.8. Special Features of Operator Classes

- [36.17. Packaging Related Objects into an Extension](#)
 - 36.17.1. Extension Files
 - 36.17.2. Extension Relocatability
 - 36.17.3. Extension Configuration Tables
 - 36.17.4. Extension Updates
 - 36.17.5. Installing Extensions Using Update Scripts
 - 36.17.6. Security Considerations for Extensions
 - 36.17.7. Extension Example
 - [36.18. Extension Building Infrastructure](#)
-

Extending PostgreSQL SQL Query Language

In the sections that follow, we will discuss how you can extend the PostgreSQL SQL query language by adding:

- functions (starting in [Section 36.3](#))
 - aggregates (starting in [Section 36.12](#))
 - data types (starting in [Section 36.13](#))
 - operators (starting in [Section 36.14](#))
 - operator classes for indexes (starting in [Section 36.16](#))
 - packages of related objects (starting in [Section 36.17](#))
-

Navigation

[Prev](#) | [Up](#) | [Next](#)

Part V. Server Programming | [Home](#) | [36.1. How Extensibility Works](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 37. Triggers

Navigation

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

Latest News

- [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)
-

Main Content

Documentation Navigation

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

[Current \(17 \)](#) / [16](#) / [15](#) / [14](#) / [13](#)

Development Versions:

[18](#) / [devel](#)

Unsupported versions:

[12](#) / [11](#) / [10](#) / [9.6](#) / [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#)

Important Notice

This documentation is for an unsupported version of PostgreSQL.
You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Chapter 37. Triggers

Chapter 37.

Triggers

[Prev](#)

[U](#) *Part V. Server*
[p](#) *Programming*

[Ho](#) *Next: trigger-*
[me](#) *definition.html*

Chapter 37. Triggers

Table of Contents

- [37.1. Overview of Trigger Behavior](#)
 - [37.2. Visibility of Data Changes](#)
 - [37.3. Writing Trigger Functions in C](#)
 - [37.4. A Complete Trigger Example](#)
-

General Information

This chapter provides general information about writing trigger functions. Trigger functions can be written in most of the available procedural languages, including *PL/pgSQL* ([Chapter 41](#)), *PL/Tcl* ([Chapter 42](#)), *PL/Perl* ([Chapter 43](#)), and *PL/Python* ([Chapter 44](#)). After reading this chapter, you should consult the chapter for your favorite procedural language to find out the language-specific details of writing a trigger in it.

It is also possible to write a trigger function in C, although most people find it easier to use one of the procedural languages. It is not currently possible to write a trigger function in the plain SQL function language.

Navigation

Prev	Up	Next
36.18. Extension Building Infrastructure	Home	37.1. Overview of Trigger Behavior

Unsupported Version Notice

This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

PostgreSQL: Documentation: 18: Chapter 38. Event Triggers

[Home](#) | [About PostgreSQL](#) | [Privacy Policy](#) | [Code of Conduct](#) | [Contact](#)
Copyright © 1996-2025 The PostgreSQL Global Development Group

Chapter 38. Event Triggers

Chapter 38. Event
Triggers

Prev Up

Part V. Server
Programming

Home Next

Prev Up
Part V. Server
Programming

Home Next

Table of Contents

- [38.1. Overview of Event Trigger Behavior](#)
 - [38.1.1. login](#)
 - [38.1.2. ddl_command_start](#)
 - [38.1.3. ddl_command_end](#)
 - [38.1.4. sql_drop](#)
 - [38.1.5. table_rewrite](#)
 - [38.1.6. Event Triggers in Aborted Transactions](#)
 - [38.1.7. Creating Event Triggers](#)
- [38.2. Writing Event Trigger Functions in C](#)
- [38.3. A Complete Event Trigger Example](#)
- [38.4. A Table Rewrite Event Trigger Example](#)

- [38.5. A Database Login Event Trigger Example](#)
-

To supplement the trigger mechanism discussed in [Chapter 37](#), PostgreSQL also provides event triggers. Unlike regular triggers, which are attached to a single table and capture only DML events, event triggers are global to a particular database and are capable of capturing DDL events.

Like regular triggers, event triggers can be written in any procedural language that includes event trigger support, or in C, but not in plain SQL.

Supporting text

To support the trigger mechanism discussed in [Chapter 37](#), PostgreSQL also provides event triggers. Unlike regular triggers, which are attached to a single table and capture only DML events, event triggers are global to a particular database and are capable of capturing DDL events.

Like regular triggers, event triggers can be written in any procedural language that includes event trigger support, or in C, but not in plain SQL.

Navigation

Prev	Up	Next
37.4. A Complete Trigger Example	Home	38.1. Overview of Event Trigger Behavior

PostgreSQL:

Documentation: 18:

Chapter 39. The Rule System

Chapter 39. The
Rule System

Prev

Up

Home

Next

[event-trigger-
database-login-
example.html](#)

[server-
programmi
ng.html](#)

[index.h
tml](#) [querytre
e.html](#)

Chapter 39. The Rule System

Table of Contents

- [39.1. The Query Tree](#)
 - [39.2. Views and the Rule System](#)
 - [39.2.1. How `SELECT` Rules Work](#)
 - [39.2.2. View Rules in Non-`SELECT` Statements](#)
 - [39.2.3. The Power of Views in PostgreSQL](#)
 - [39.2.4. Updating a View](#)
 - [39.3. Materialized Views](#)
 - [39.4. Rules on `INSERT`, `UPDATE`, and `DELETE`](#)
 - [39.4.1. How Update Rules Work](#)
 - [39.4.2. Cooperation with Views](#)
 - [39.5. Rules and Privileges](#)
 - [39.6. Rules and Command Status](#)
 - [39.7. Rules Versus Triggers](#)
-

This chapter discusses the rule system in *PostgreSQL*. Production rule systems are conceptually simple, but there are many subtle points involved in actually using them.

Some other database systems define active database rules, which are usually stored procedures and triggers. In *PostgreSQL*, these can be implemented using functions and triggers as well.

The rule system (more precisely speaking, the query rewrite rule system) is totally different from stored procedures and triggers. It modifies queries to take rules into consideration, and then passes the modified query to the query planner for planning and execution. It is very powerful, and can be used for many things such as query language procedures, views, and versions. The theoretical foundations and the power of this rule system are also discussed in [ston90b](#) and [ong90](#).

Prev	Up	Next
event-trigger-database-login-example.html	server-programming.html	querytree.html

38.5. A Database Login Event Trigger Example

PostgreSQL documentation for chapter 39. The Rule System.

Copyright © 1996-2025 The PostgreSQL Global Development Group

PostgreSQL: Documentation: 18: Chapter 40. Procedural Languages

Navigation

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

News

May 8, 2025: [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#) | [PostgreSQL 18 Beta 1 Released!](#)

Documentation for PostgreSQL 18

Supported Versions

- [Current](#) (17)
- [16](#)
- [15](#)
- [14](#)
- [13](#)

Development Versions

- [18](#)
- [devel](#)

Unsupported Versions

- [12](#)
- [11](#)
- [10](#)
- [9.6](#)
- [9.5](#)
- [9.4](#)
- [9.3](#)
- [9.2](#)
- [9.1](#)
- [9.0](#)
- [8.4](#)
- [8.3](#)
- [8.2](#)
- [8.1](#)
- [8.0](#)
- [7.4](#)
- [7.3](#)
- [7.2](#)
- [7.1](#)

Note: This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of

the other supported versions listed above instead.

Chapter 40. Procedural Languages

Prev	Up	Part V. Server Programming	Home	Next
Rules & Triggers	Server Programming	Part V. Server Programming	Home	Installing Procedural Languages

Chapter 40. Procedural Languages

Table of Contents

- [40.1. Installing Procedural Languages](#)
-

PostgreSQL allows user-defined functions to be written in languages besides SQL and C. These are called *procedural languages* (PLs). For a function written in a procedural language, the database server relies on a handler that understands the language's source text. This handler could perform parsing, syntax analysis, execution, or serve as *glue* between PostgreSQL and an existing language implementation. Handlers are C language functions compiled into shared objects loaded on demand.

Currently, PostgreSQL supports four procedural languages in its standard distribution: PL/pgSQL ([Chapter 41](#)), PL/Tcl ([Chapter 42](#)), PL/Perl ([Chapter 43](#)), and PL/Python ([Chapter 44](#)). Additional languages can be found in [Appendix H](#), and users can define their own, with guidance in [Chapter 57](#).

Navigation

- [Prev](#): 39.7. Rules Versus Triggers
 - [Up](#)
 - [Next](#): 40.1. Installing Procedural Languages
-

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

© 1996-2025 The PostgreSQL Global Development Group

Navigation

- [Home](#)
 - [About](#)
 - [Download](#)
 - [Documentation](#)
 - [Community](#)
 - [Developers](#)
 - [Support](#)
 - [Donate](#)
 - [Your account](#)
-

News

- [PostgreSQL 17.5, 16.9, 15.13, 14.18, and 13.21 Released!](#)
 - [PostgreSQL 18 Beta 1 Released!](#)
-

Main Content

Chapter 41. PL/pgSQL — SQL Procedural Language

[Documentation](#) → [PostgreSQL 18](#)

Supported Versions:

- [Current](#) (17)
- [16](#), [15](#), [14](#), [13](#), [12](#), [11](#), [10](#), [9.6](#), [9.5](#), [9.4](#), [9.3](#), [9.2](#), [9.1](#), [9.0](#), [8.4](#), [8.3](#), [8.2](#), [8.1](#), [8.0](#)

Development Versions:

- [18](#), [devel](#)

Unsupported Versions:

- [12](#), [11](#), [10](#), [9.6](#), [9.5](#), [9.4](#), [9.3](#), [9.2](#), [9.1](#), [9.0](#), [8.4](#), [8.3](#), [8.2](#), [8.1](#), [8.0](#), [7.4](#), [7.3](#), [7.2](#), [7.1](#)

Note: This documentation is for an unsupported version of PostgreSQL. You may want to view the same page for the [current](#) version, or one of the other supported versions listed above instead.

Navigation Table

Prev	Up	Next
xplang-install.html	server-programming.html	index.html
40.1. Installing Procedural Languages		41.1. Overview

Chapter Title

Chapter 41. PL/pgSQL — SQL Procedural Language

Table of Contents

- [41.1. Overview](#)
 - [41.1.1. Advantages of Using PL/pgSQL](#)
 - [41.1.2. Supported Argument and Result Data Types](#)
- [41.2. Structure of PL/pgSQL](#)
- [41.3. Declarations](#)
 - [41.3.1. Declaring Function Parameters](#)
 - [41.3.2. ALIAS](#)
 - [41.3.3. Copying Types](#)
 - [41.3.4. Row Types](#)
 - [41.3.5. Record Types](#)
 - [41.3.6. Collation of PL/pgSQL Variables](#)

- 41.4. Expressions
- 41.5. Basic Statements
 - 41.5.1. Assignment
 - 41.5.2. Executing SQL Commands
 - 41.5.3. Executing a Command with a Single-Row Result
 - 41.5.4. Executing Dynamic Commands
 - 41.5.5. Obtaining the Result Status
 - 41.5.6. Doing Nothing At All
- 41.6. Control Structures
 - 41.6.1. Returning from a Function
 - 41.6.2. Returning from a Procedure
 - 41.6.3. Calling a Procedure
 - 41.6.4. Conditionals
 - 41.6.5. Simple Loops
 - 41.6.6. Looping through Query Results
 - 41.6.7. Looping through Arrays
 - 41.6.8. Trapping Errors
 - 41.6.9. Obtaining Execution Location Information
- 41.7. Cursors
 - 41.7.1. Declaring Cursor Variables
 - 41.7.2. Opening Cursors
 - 41.7.3. Using Cursors
 - 41.7.4. Looping through a Cursor's Result
- 41.8. Transaction Management
- 41.9. Errors and Messages
 - 41.9.1. Reporting Errors and Messages
 - 41.9.2. Checking Assertions
- 41.10. Trigger Functions
 - 41.10.1. Triggers on Data Changes
 - 41.10.2. Triggers on Events
- 41.11. PL/pgSQL under the Hood
 - 41.11.1. Variable Substitution
 - 41.11.2. Plan Caching
- 41.12. Tips for Developing in PL/pgSQL
 - 41.12.1. Handling of Quotation Marks
 - 41.12.2. Additional Compile-Time and Run-Time Checks
- 41.13. Porting from Oracle PL/SQL
 - 41.13.1. Porting Examples

- [41.13.2. Other Things to Watch For](#)
 - [41.13.3. Appendix](#)
-

Support Links

[Prev: Installing Procedural Languages](#) | [Up: Server Programming](#) | [Next: Chapter 41. Overview](#)

Footer

[Privacy Policy](#) | [Code of Conduct](#) | [About PostgreSQL](#) | [Contact](#)

Copyright © 1996-2025 The PostgreSQL Global Development Group