

PostgreSQL

Find the detailed version of this roadmap along with resources and other roadmaps
<https://roadmap.sh>

Introduction

- What are Relational Databases?
- RDBMS Benefits and Limitations
- PostgreSQL vs Other RDBMS
- PostgreSQL vs NoSQL Databases

Basic RDBMS Concepts

- Queries
- Data Types
- Rows
- Columns
- Tables
- Schemas
- Databases
- Object Model

- High Level Database Concepts
 - ACID
 - MVCC
 - Transactions
- Write-ahead Log
- Query Processing

- Relational Model
 - Domains
 - Attributes
 - Tuples
 - Relations
 - Constraints
 - NULL

Installation and Setup

- Using Docker
- Package Managers

- Managing Postgres
 - Using `systemd`
 - Using `pg_ctl`
 - Using `pg_ctlcluster`

- Connect using `psql`
- Deployment in Cloud

Learn SQL Concepts

- DDL Queries
 - For Schemas
 - For Tables
 - Data Types
- DML Queries
 - Querying Data
 - Filtering Data
 - Modifying Data
 - Joining Tables

- Import / Export using `COPY`
- Advanced Topics
 - Transactions
 - CTE
 - Subqueries
 - Lateral Join
 - Grouping
 - Set Operations

Configuring PostgreSQL

- Resources Usage
- Write-ahead Log
- Vacuums
- Replication
- Query Planner
- Checkpoints
- Background Writer

Following postgres.conf configurations and more

Adding Extra Extensions

- Default Privileges
- Grant / Revoke
- Object Privileges

Postgres Security Concepts

- Authentication Models
- Roles
- pg_hba.conf
- SSL Settings

- Advanced Topics
- Row-Level Security
- SELinux

- Logical Replication
- Streaming Replication
- Replication

- PgBouncer
- PgBouncer Alternatives
- Connection Pooling

Develop Infrastructure Skills

- Kubernetes Deployment
 - Simple Stateful Setup
 - Helm
 - Operators
- Cluster Management
 - Patroni
 - Patroni Alternatives
- Upgrade Procedures
 - Using `pg_upgrade`
 - Using Logical Replication
- Monitoring
 - Prometheus
 - Zabbix
 - check_pgactivity
 - temBoard
 - check_pgbackrest
- Anonymization
 - PostgreSQL Anonymizer
- Load Balancing & Service Discovery
 - HAProxy
 - Consul
 - KeepAlived
 - Etcad
- Resource Usage and Provisioning, Capacity Planning

Backup / Recovery Tools

- 3rd Party
 - barman
 - WAL-G
 - pgbackrest
 - pg_probackup
- Built-in Tools
 - pg_dump
 - pg_dumpall
 - pg_restore
 - pg_basebackup
- Backup Validation Procedures

Learn to Automate Routines

- Shell Scripts
- Any Programming Language
- Terraform
- Configuration Management
 - Ansible
 - Salt
 - Chef
 - Puppet

Application Skills

- Migrations
 - Practical Patterns and Antipatterns
 - liquibase, sqitch, Bytebase, ora2pg etc
 - Migration Tools
- Bulk Loading and Processing Data
- Data Partitioning
- Sharding Patterns
- Data Normalization / Normal Forms
- Queues
 - Practical Patterns and Antipatterns
 - Skytools PGQ

Advanced Topics

- Low Level Internals
 - Processes and memory architecture
 - Vacuum Processing
 - Buffer Management
 - Lock Management
 - Physical Storage and File Layout
 - System Catalog
- Fine Grained Tuning
 - Per-user, Per-Database Settings
 - Storage Parameters
 - Workload-dependant tuning: OLTP, OLAP, HTAP
- Advanced SQL
 - PL/pgSQL
 - Procedures and Functions
 - Triggers
 - Recursive CTE
 - Aggregate and window functions

Troubleshooting Techniques

- Postgres System Views
 - pg_stat_activity
 - pg_stat_statements
- Postgres Tools
 - pgcenter
- Query Analysis
 - EXPLAIN
 - Depesz
 - PEV2
 - Tensor
 - explain.dalibo.com
- Troubleshooting Methods
 - USE
 - RED
 - Golden Signals
- Operating System Tools
 - top
 - sysstat
 - iotop
- Profiling Tools
 - gdb
 - strace
 - ebpf
 - perf-tools
 - Core Dumps
- Log Analysis
 - pgBadger
 - pgCluu
 - awk
 - grep
 - sed

SQL Optimization Techniques

- Indexes and their Usecases
 - B-Tree
 - Hash
 - GiST
 - SP-GiST
 - GIN
 - BRIN

- Schema Design Patterns / Anti-patterns
- SQL Query Patterns / Anti-patterns

Get Involved in Development

- Mailing Lists
- Reviewing Patches
- Writing Patches

Keep Learning