

# Table of Contents

<b>Preface</b>	1
<b>Chapter 1: PostgreSQL 13 Overview</b>	7
<b>What's new in PostgreSQL 13?</b>	7
Digging into SQL and developer-related topics	8
Improving psql command-line handling	8
Improving pgbench	8
Generating random UUIDs more easily	9
Dropping databases faster	9
Adding ALTER TABLE ... DROP EXPRESSION ...	10
Making use of performance improvements	11
Deduplication of B-tree indexes	11
Adding incremental sorting	12
Adding -j 8 to reindexdb	14
Allowing hash aggregates to spill to disk	15
Speeding up PL/pgSQL	17
Parallelizing VACUUM operations	17
Allowing skipping of WAL for full table writes	18
Additional performance improvements	18
Making monitoring more powerful	19
Additional system views	19
<b>Summary</b>	21
<b>Chapter 2: Understanding Transactions and Locking</b>	23
<b>Working with PostgreSQL transactions</b>	24
Handling errors inside a transaction	27
Making use of SAVEPOINT	28
Transactional DDLs	29
<b>Understanding basic locking</b>	30
Avoiding typical mistakes and explicit locking	32
Considering alternative solutions	34
<b>Making use of FOR SHARE and FOR UPDATE</b>	35
<b>Understanding transaction isolation levels</b>	39
Considering Serializable Snapshot Isolation transactions	40
<b>Observing deadlocks and similar issues</b>	41
<b>Utilizing advisory locks</b>	43
<b>Optimizing storage and managing cleanup</b>	44
Configuring VACUUM and autovacuum	45
Digging into transaction wraparound-related issues	46
A word on VACUUM FULL	48
Watching VACUUM at work	48
Limiting transactions by making use of snapshot too old	51

Making use of more VACUUM features	52
<b>Summary</b>	53
<b>Questions</b>	53
<b>Chapter 3: Making Use of Indexes</b>	55
<b>Understanding simple queries and the cost model</b>	56
Making use of EXPLAIN	57
Digging into the PostgreSQL cost model	59
Deploying simple indexes	61
Making use of sorted output	62
Using more than one index at a time	63
Using bitmap scans effectively	64
Using indexes in an intelligent way	65
Understanding index de-duplication	67
<b>Improving speed using clustered tables</b>	68
Clustering tables	71
Making use of index-only scans	72
<b>Understanding additional B-tree features</b>	73
Combined indexes	74
Adding functional indexes	75
Reducing space consumption	76
Adding data while indexing	77
<b>Introducing operator classes</b>	78
Creating an operator class for a B-tree	80
Creating new operators	80
Creating operator classes	84
Testing custom operator classes	84
<b>Understanding PostgreSQL index types</b>	85
Hash indexes	86
GiST indexes	86
Understanding how GiST works	87
Extending GiST	88
GIN indexes	89
Extending GIN	90
SP-GiST indexes	91
BRIN indexes	92
Extending BRIN indexes	93
Adding additional indexes	93
<b>Achieving better answers with fuzzy searching</b>	95
Taking advantage of pg_trgm	95
Speeding up LIKE queries	97
Handling regular expressions	98
<b>Understanding full-text search</b>	99
Comparing strings	100
Defining GIN indexes	101
Debugging your search	102

Gathering word statistics	103
Taking advantage of exclusion operators	103
<b>Summary</b>	104
<b>Questions</b>	105
<b>Chapter 4: Handling Advanced SQL</b>	107
<b>Introducing grouping sets</b>	108
Loading some sample data	108
Applying grouping sets	109
Investigating performance	111
Combining grouping sets with the FILTER clause	113
<b>Making use of ordered sets</b>	113
<b>Understanding hypothetical aggregates</b>	116
<b>Utilizing windowing functions and analytics</b>	117
Partitioning data	118
Ordering data inside a window	119
Using sliding windows	120
Understanding the subtle difference between ROWS and RANGE	123
Removing duplicates using EXCLUDE TIES and EXCLUDE GROUP	125
Abstracting window clauses	127
Using on-board windowing functions	127
The rank and dense_rank functions	128
The ntile() function	129
The lead() and lag() functions	130
The first_value(), nth_value(), and last_value() functions	133
The row_number() function	135
<b>Writing your own aggregates</b>	135
Creating simple aggregates	136
Adding support for parallel queries	139
Improving efficiency	140
Writing hypothetical aggregates	142
<b>Summary</b>	144
<b>Chapter 5: Log Files and System Statistics</b>	145
<b>Gathering runtime statistics</b>	145
Working with PostgreSQL system views	146
Checking live traffic	146
Inspecting databases	149
Inspecting tables	152
Making sense of pg_stat_user_tables	153
Digging into indexes	155
Tracking the background worker	156
Tracking, archiving, and streaming	157
Checking SSL connections	160
Inspecting transactions in real time	161
Tracking VACUUM and CREATE INDEX progress	162
Using pg_stat_statements	164

<b>Creating log files</b>	169
Configuring the postgresql.conf file	169
Defining log destination and rotation	170
Configuring syslog	171
Logging slow queries	172
Defining what and how to log	172
<b>Summary</b>	175
<b>Questions</b>	176
<b>Chapter 6: Optimizing Queries for Good Performance</b>	177
<b>Learning what the optimizer does</b>	177
A practical example – How the query optimizer handles a sample query	178
Evaluating join options	178
Nested loops	179
Hash joins	179
Merge joins	179
Applying transformations	180
Applying equality constraints	181
Exhaustive searching	182
Checking out execution plans	182
Making the process fail	184
Constant folding	185
Understanding function inlining	186
Introducing join pruning	188
Speedup set operations	189
<b>Understanding execution plans</b>	191
Approaching plans systematically	191
Making EXPLAIN more verbose	193
Spotting problems	194
Spotting changes in runtime	194
Inspecting estimates	195
Inspecting buffer usage	197
Fixing high buffer usage	198
<b>Understanding and fixing joins</b>	199
Getting joins right	199
Processing outer joins	201
Understanding the join_collapse_limit variable	202
<b>Enabling and disabling optimizer settings</b>	203
Understanding genetic query optimization	206
<b>Partitioning data</b>	208
Creating inherited tables	208
Applying table constraints	210
Modifying inherited structures	212
Moving tables in and out of partitioned structures	213
Cleaning up data	214
Understanding PostgreSQL 13.x partitioning	215
<b>Adjusting parameters for good query performance</b>	217

Speeding up sorting	220
Speeding up administrative tasks	223
<b>Making use of parallel queries</b>	224
What is PostgreSQL able to do in parallel?	228
Parallelism in practice	229
<b>Introducing JIT compilation</b>	230
Configuring JIT	231
Running queries	231
<b>Summary</b>	234
<b>Chapter 7: Writing Stored Procedures</b>	235
<b>Understanding stored procedure languages</b>	235
Understanding the fundamentals of stored procedures versus functions	237
The anatomy of a function	238
Introducing dollar quoting	239
Making use of anonymous code blocks	240
Using functions and transactions	240
<b>Exploring various stored procedure languages</b>	242
Introducing PL/pgSQL	243
Handling quoting and the string format	243
Managing scopes	247
Understanding advanced error handling	248
Making use of GET DIAGNOSTICS	250
Using cursors to fetch data in chunks	251
Utilizing composite types	255
Writing triggers in PL/pgSQL	256
Writing stored procedures in PL/pgSQL	261
Introducing PL/Perl	263
Using PL/Perl for data type abstraction	264
Deciding between PL/Perl and PL/PerlU	265
Making use of the SPI interface	266
Using the SPI for set returning functions	267
Escaping in PL/Perl and support functions	269
Sharing data across function calls	269
Writing triggers in Perl	270
Introducing PL/Python	270
Writing simple PL/Python code	271
Using the SPI interface	271
Handling errors	273
<b>Improving functions</b>	274
Reducing the number of function calls	275
Using cached plans	276
Assigning costs to functions	276
<b>Using functions for various purposes</b>	277
<b>Summary</b>	279
<b>Questions</b>	279
<b>Chapter 8: Managing PostgreSQL Security</b>	281

<b>Managing network security</b>	281
Understanding bind addresses and connections	282
Inspecting connections and performance	284
Living in a world without TCP	284
Managing pg_hba.conf	286
Handling SSL	289
Handling instance-level security	291
Creating and modifying users	294
Defining database-level security	297
Adjusting schema-level permissions	299
Working with tables	302
Handling column-level security	303
Configuring default privileges	304
<b>Digging into RLS</b>	305
<b>Inspecting permissions</b>	310
<b>Reassigning objects and dropping users</b>	313
<b>Summary</b>	314
<b>Questions</b>	315
<b>Chapter 9: Handling Backup and Recovery</b>	317
<b>Performing simple dumps</b>	317
Running pg_dump	318
Passing passwords and connection information	319
Using environment variables	320
Making use of .pgpass	321
Using service files	321
Extracting subsets of data	322
<b>Handling various formats</b>	323
<b>Replaying backups</b>	325
<b>Handling global data</b>	326
<b>Summary</b>	327
<b>Questions</b>	327
<b>Chapter 10: Making Sense of Backups and Replication</b>	329
<b>Understanding the transaction log</b>	330
Looking at the transaction log	331
Understanding checkpoints	332
Optimizing the transaction log	332
<b>Transaction log archiving and recovery</b>	334
Configuring for archiving	334
Configuring the pg_hba.conf file	335
Creating base backups	336
Reducing the bandwidth of backups	337
Mapping tablespaces	338
Using different formats	338
Testing transaction log archiving	339

Replaying the transaction log	340
Finding the right timestamp	342
Cleaning up the transaction log archive	344
<b>Setting up asynchronous replication</b>	346
Performing a basic setup	346
Improving security	348
Halting and resuming replication	349
Checking replication to ensure availability	350
Performing failovers and understanding timelines	353
Managing conflicts	354
Making replication more reliable	356
<b>Upgrading to synchronous replication</b>	357
Adjusting durability	359
<b>Making use of replication slots</b>	361
Handling physical replication slots	362
Handling logical replication slots	364
Use cases for logical replication slots	366
<b>Making use of the CREATE PUBLICATION and CREATE SUBSCRIPTION commands</b>	367
<b>Summary</b>	369
<b>Questions</b>	369
<b>Chapter 11: Deciding on Useful Extensions</b>	371
<b>Understanding how extensions work</b>	372
Checking for available extensions	373
<b>Making use of contrib modules</b>	376
Using the adminpack module	376
Applying bloom filters	378
Deploying btree_gist and btree_gin	380
dblink – considering phasing out	382
Fetching files with file_fdw	382
Inspecting storage using pageinspect	384
Investigating caching with pg_buffercache	386
Encrypting data with pgcrypto	388
Prewarming caches with pg_prewarm	388
Inspecting performance with pg_stat_statements	390
Inspecting storage with pgstattuple	390
Fuzzy searching with pg_trgm	392
Connecting to remote servers using postgres_fdw	392
Handling mistakes and typos	396
<b>Other useful extensions</b>	397
<b>Summary</b>	398
<b>Chapter 12: Troubleshooting PostgreSQL</b>	399
<b>Approaching an unknown database</b>	400
<b>Inspecting pg_stat_activity</b>	400

Querying pg_stat_activity	401
Treating Hibernate statements	402
Figuring out where queries come from	403
<b>Checking for slow queries</b>	403
Inspecting individual queries	405
Digging deeper with perf	405
<b>Inspecting the log</b>	407
<b>Checking for missing indexes</b>	407
<b>Checking for memory and I/O</b>	408
<b>Understanding noteworthy error scenarios</b>	410
Facing clog corruption	411
Understanding checkpoint messages	412
Managing corrupted data pages	413
Careless connection management	414
Fighting table bloat	414
<b>Summary</b>	415
<b>Questions</b>	415
<b>Chapter 13: Migrating to PostgreSQL</b>	417
<b>Migrating SQL statements to PostgreSQL</b>	417
Using lateral joins	418
Supporting lateral joins	418
Using grouping sets	418
Supporting grouping sets	419
Using the WITH clause – common table expressions	419
Supporting the WITH clause	420
Using the WITH RECURSIVE clause	421
Supporting the WITH RECURSIVE clause	421
Using the FILTER clause	421
Supporting the FILTER clause	422
Using windowing functions	422
Supporting windowing and analytics	422
Using ordered sets – the WITHIN GROUP clause	423
Supporting the WITHIN GROUP clause	423
Using the TABLESAMPLE clause	424
Supporting the TABLESAMPLE clause	424
Using limit/offset	425
Supporting the FETCH FIRST clause	425
Using the OFFSET clause	425
Supporting the OFFSET clause	426
Using temporal tables	426
Supporting temporal tables	426
Matching patterns in time series	427
<b>Moving from Oracle to PostgreSQL</b>	427
Using the oracle_fdw extension to move data	427
Using ora_migrator for fast migration	429

CYBERTEC Migrator – migration for the "big boys"	431
Using Ora2Pg to migrate from Oracle	431
Common pitfalls	434
Handling data in MySQL and MariaDB	436
Changing column definitions	437
Handling null values	439
Expecting problems	440
Migrating data and schemas	440
Using pg_chameleon	441
Using FDWs	442
<b>Summary</b>	444
<b>Other Books You May Enjoy</b>	445
<b>Index</b>	449

---