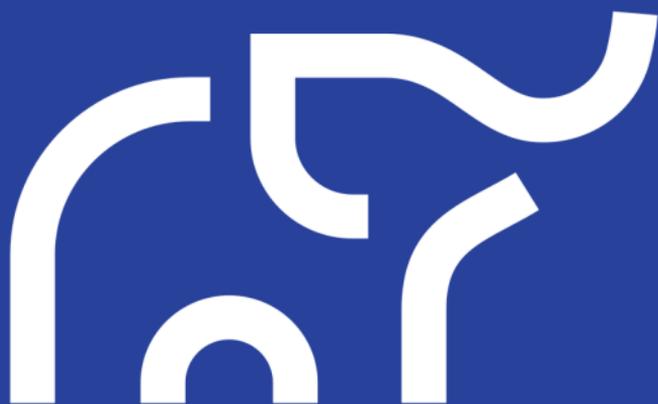


 PostgresPro

# Postgres

P. Luzanov, E. Rogov, I. Levshin  
Translated by L. Mantrova, A. Meleshko



The first experience

17

# Introduction

We have written this small book for those who only start getting acquainted with the world of PostgreSQL. From this book, you will learn:

I	PostgreSQL – what is it all about? .....	3
II	What's new in PostgreSQL 17 .....	15
III	Installation on Linux and Windows .....	25
IV	Connecting to a server, writing SQL queries, and using transactions .....	35
V	Learning the SQL language on a demo database ....	61
VI	Using PostgreSQL with your application.....	89
VII	Minimal server setup .....	103
VIII	About a useful pgAdmin application .....	111
IX	Advanced features: full-text search, .....	117
	JSON format, .....	124
	foreign data wrappers .....	137
X	Education and certification opportunities .....	149
XI	Keeping up with all updates .....	177
XII	About the Postgres Professional company .....	181

We hope that our book will make your first experience with PostgreSQL more pleasant and help you blend into the PostgreSQL community. Good luck!



# I About PostgreSQL

PostgreSQL is the most feature-rich free and open-source database system. Originally developed in an academic environment, it has successfully united a large community of developers over the years. Nowadays, PostgreSQL offers everything that most customers need, and it is actively used all over the world to create high-load business-critical systems.

## Some History

Modern PostgreSQL originates from the POSTGRES project, which was led by Michael Stonebraker, professor of the University of California, Berkeley. Before this work, Michael Stonebraker managed the development of INGRES, one of the first relational database systems. POSTGRES emerged as an effort to rethink previous work and address the limitations of INGRES's rigid type system.

The project was started in 1985, and by 1988 a number of scientific articles had been published that described the data model, POSTQUEL query language (SQL was not an accepted standard at the time), and data storage structure.

- 4 POSTGRES is sometimes considered to be a so-called post-  
i relational database system. The relational model had long been criticized for its restrictions, which were an inevitable trade-off for its strictness and simplicity. As computer technologies were spreading in all spheres of life, new types of applications started to appear, and databases had to support custom data types and such features as inheritance or creating and managing complex objects.

The first version of this database system appeared in 1989. It was being improved and enhanced for several years, but in 1993, when version 4.2 was released, the project was shut down. However, despite its official cancellation, UC Berkeley alumni Andrew Yu and Jolly Chen revived the project and resumed its development in 1994, taking advantage of its liberal BSD license and open source. They replaced POSTQUEL query language with SQL, which had become a generally accepted standard by that time. The project was renamed to Postgres95.

In 1996, it became obvious that the Postgres95 name would not stand the test of time, and a new name was selected: PostgreSQL. This name reflects the connection both with the original POSTGRES project and the SQL adoption. This name may be quite hard to articulate, but nevertheless, we should pronounce it as “Post-Gres-Q-L” or simply “postgres,” but not as “postgre.”

The first PostgreSQL release had version 6.0, keeping the original numbering scheme. The project grew, and its management was taken over by at first a small group of active users and developers, which was named PostgreSQL Global Development Group.

## Development

5  
i

All the main decisions about developing and releasing new PostgreSQL versions are taken by the Core team, which consists of seven people at the moment.

Apart from occasional contributors, there is a core group of developers who make significant contributions to PostgreSQL. They are called major contributors. There is also a group of committers, who have the write access to the source code repository. The group's membership evolves over time, with new developers joining the community while others move on. The current list of developers is published on PostgreSQL's official website: [postgresql.org/community/contributors](https://www.postgresql.org/community/contributors).

The contribution of Russian developers into PostgreSQL is compelling. This is arguably the largest global open-source project with such a vast Russian representation.

Vadim Mikheev, a software programmer from Krasnoyarsk who used to be a member of the Core team, played an important role in PostgreSQL evolution and development. He created such key core features as multi-version concurrency control (MVCC), vacuum, write-ahead log (WAL), subqueries, triggers. Vadim is not involved with the project anymore.

In 2015, Oleg Bartunov, a professional astronomer and research scientist at Sternberg Astronomical Institute of Lomonosov Moscow State University, teamed up with Teodor Sigaev and Alexander Korotkov to start the Postgres Professional company, which is now the main talent foundry in Russia when it comes to database system development.

The main areas of their contribution are PostgreSQL localization (national encodings and Unicode support), full-text

- 6 search, working with arrays and semi-structured data (hstore,  
i json, jsonb), new index methods (GiST, SP-GiST, GIN and RUM,  
Bloom). They have also created a lot of popular extensions.

The PostgreSQL release cycle usually takes about a year. In this timeframe, the community receives patches with bug fixes, updates, and new features from contributors worldwide. Traditionally, all patches are discussed in the `pgsql-hackers` mailing list. If the community finds the idea useful, its implementation is correct, and the code passes a mandatory code review by other developers, the patch is included into the next release.

At some point (usually in spring, about half a year before the release), code stabilization is announced: all new features get postponed until the next version, only bug fixes and improvements for the already included patches are accepted. Within the release cycle, beta versions appear. Closer to the end of the release cycle a release candidate is built, and soon a new major version of PostgreSQL is released.

Major versions used to be defined by two numbers, but in 2017 it was decided to start using a single number. Thus, version 9.6 was followed by PostgreSQL 10, while the latest available version is PostgreSQL 17, which was released in September 2024.

As the new version is being prepared, developers can find and fix bugs in it. The most critical fixes are backported to the previous versions. The community usually releases updates quarterly; these minor versions accumulate such fixes. For example, version 12.5 contains bug fixes for the 12.4 release, while version 17.1 provides fixes for PostgreSQL 17.0.

## Support

PostgreSQL Global Development Group supports major releases for five years. Both support and development are managed through mailing lists. A properly filed bug report has a high likelihood of being addressed quickly; bug fixes can be released in as little as 24 hours.

Aside from the community support, 24x7 commercial support for PostgreSQL is also provided by a number of companies in different countries, including Postgres Professional ([www.postgrespro.com](http://www.postgrespro.com)).

## Current State

PostgreSQL is one of the most popular databases. Based on the solid foundation of academic development, over several decades PostgreSQL has evolved into an enterprise-level product that is now a real alternative to commercial databases. You can see it for yourself by looking at the key features of PostgreSQL 17, which is the latest released version right now.

## Reliability and Stability

Reliability is especially important in enterprise-level applications that handle business-critical data. For this purpose, PostgreSQL provides support for hot standby servers, point-in-time recovery, different types of replication (synchronous, asynchronous, cascade).

## Security

PostgreSQL supports secure SSL connections and provides various authentication methods, such as password authentication (including SCRAM), client certificates, and external authentication services (LDAP, RADIUS, PAM, Kerberos).

For user management and database access control, the following features are provided:

- creating and managing new users and group roles
- role- and group-based access control to database objects
- row-level and column-level security
- SELinux support via a built-in SE-PostgreSQL functionality (Mandatory Access Control)

Russian Federal Service for Technical and Export Control has certified a custom PostgreSQL version released by Postgres Professional for use in data processing systems for personal data and classified information.

## Conformance to the SQL Standard

As the ANSI SQL standard is evolving, its support is constantly being added to PostgreSQL. This is true for all versions of the standard, from SQL-92 to the most recent SQL:2023. For example, PostgreSQL 17 fully supports all JSON operations as described in the SQL:2016 Standard.

In general, PostgreSQL provides a high rate of conformance to the SQL standard, supporting 170 out of 177 mandatory features and many optional ones.

## Transaction Support

PostgreSQL provides full support for ACID properties and efficient transaction isolation based on the multi-version concurrency control (MVCC). This method avoids locking in all cases except for concurrent updates to the same row by different processes. Reading transactions never block writing ones, and writing never blocks reading.

This is true even for the serializable isolation level, which is the strictest one. Using an innovative Serializable Snapshot Isolation system, this level ensures that there are no serialization anomalies and guarantees that concurrent transaction execution produces the same result as sequential one.

## For Application Developers

Application developers get a rich toolset for creating applications of any type:

- Support for various server programming languages: built-in PL/pgSQL (which is closely integrated with SQL), C for performance-critical tasks, Perl, Python, Tcl, as well as JavaScript, Java, etc.
- APIs to access the database from applications written in virtually any language, including the standard ODBC and JDBC APIs.
- A rich set of database objects that allow you to efficiently implement the logic of any complexity on the server side: tables and indexes, sequences, integrity constraints, views and materialized views, partitioning, subqueries and WITH-queries (including recursive ones), aggregate and window functions, stored functions, triggers, etc.

- A flexible full-text search system that supports a variety of languages, extended with efficient index access methods.
- Semi-structured data typical of NoSQL: hstore (storage of key–value pairs), xml, json (represented as text or in a more robust jsonb binary format).
- Foreign Data Wrappers. This feature allows adding new data sources as external tables by the SQL/MED standard. You can use any major database as an external data source. PostgreSQL provides full support for foreign data, including write access and distributed query execution.

## Scalability and Performance

PostgreSQL takes advantage of the modern multi-core CPU architecture. Its performance grows almost linearly as the number of cores increases.

PostgreSQL can parallelize queries and some commands (such as index creation and vacuuming). In this mode, reads and joins are performed by several concurrent processes. JIT-compilation of queries can speed up operations thanks to better use of hardware resources. Each PostgreSQL version adds new parallelization features.

Horizontal scaling can rely on both physical and logical replication. It allows you to build PostgreSQL-based clusters to achieve high performance, fault tolerance, and geo-distribution. Some examples of such systems are Citus (Citus-data), Postgres-BDR (2ndQuadrant), Multimaster and BIHA (Postgres Professional), Patroni (Zalando).

PostgreSQL relies on a cost-based query planner. Using the collected statistics and taking into account both disk operations and CPU time in its mathematical models, the planner can optimize even the most complex queries. It can use all access paths and join methods available in state-of-the-art commercial database systems.

## Indexing

PostgreSQL provides various types of indexes. Apart from traditional B-trees, you can use many other access methods.

- Hash, a hash-based index. Unlike B-trees, such indexes work only for equality checks, but in some cases they can prove to be more efficient and compact.
- GiST, a generalized balanced search tree. This access method is used for the data that cannot be ordered. For example, R-trees index points on a plane and facilitate fast k-nearest neighbor (k-NN) searches or indexing overlapping intervals.
- SP-GiST, a generalized non-balanced tree based on dividing the search space into non-intersecting nested partitions. For example, quad-trees for spatial data and radix trees for text strings.
- GIN, a generalized inverted index used for compound multi-element values. It is mainly applied in full-text search to

12  
i

find documents that contain the words used in the search query. Another example is search for elements in data arrays.

- RUM,  
an enhancement of the GIN method for full-text search. Available as an extension, this index type can speed up phrase search and return the results in the order of relevance without any additional computations.
- BRIN,  
a compact structure that provides a trade-off between the index size and search efficiency. Such index is useful for huge clustered tables.
- Bloom,  
an index based on the Bloom filter. Having a compact representation, this index can quickly filter out non-matching tuples, but the remaining ones have to be re-checked.

Many index types can be built upon both a single column and multiple columns. Regardless of the type, you can build indexes not only on columns, but also on arbitrary expressions. It is also possible to create partial indexes for specific sets of rows. Covering indexes can speed up queries, since all the required data is retrieved from the index itself, avoiding heap access.

The planner can use a bitmap scan, which allows combining several indexes together for faster access.

## **Cross-Platform Support**

PostgreSQL runs both on Unix operating systems (including server and client Linux distributions, FreeBSD, Solaris, and macOS) and on Windows systems.

Its portable open-source C code allows building PostgreSQL on a variety of platforms, even if there is no package supported by the community. 13  
i

## Extensibility

One of the main advantages of PostgreSQL architecture is extensibility. Without changing the core system code, users can add various features, such as:

- data types
- functions and operators to support new data types
- index and table access methods
- server programming languages
- foreign data wrappers
- loadable extensions

Comprehensive extension support allows for the development of new features of any complexity, which can be installed on demand without modifying the PostgreSQL core. For example, the following complex systems are built as extensions:

- CitusDB, which implements massively parallel query execution and data distribution between different PostgreSQL instances (sharding).
- PostGIS, one of the most popular and powerful geoinformation data processing systems.

- 14 • TimescaleDB,  
i which provides support for time-series data, including special partitioning and sharding.

The standard PostgreSQL 17 distribution alone includes about fifty extensions that have proved to be useful and reliable.

## Availability

A liberal PostgreSQL license, which is similar to BSD and MIT licenses, allows unrestricted use of PostgreSQL; you may also modify PostgreSQL code without any limitations and integrate it into other products, including commercial and closed-source software.

## Independence

PostgreSQL does not belong to any company; it is developed by the international community, which includes developers from all over the world. It means that systems using PostgreSQL do not depend on a particular vendor, thus keeping the investment safe in any circumstances.

# II What's New in PostgreSQL 17

If you are familiar with the previous versions of PostgreSQL, this chapter can give you a sense of what has changed over the past year. It mentions only some of the updates; for the full list of changes, see the Release Notes: [postgresql.org/docs/postgresql-17/release-17](https://www.postgresql.org/docs/postgresql-17/release-17). You can also follow our blog to keep up with the news: <https://habr.com/en/companies/postgrespro/articles/>.

## SQL Commands

**Improved MERGE command.** The new `WHEN NOT MATCHED BY SOURCE` clause now lets you work with rows from the target relation which are not in the source, the `RETURNING` clause is now supported, and you can use views as targets same way you did with tables.

**The COPY FROM command can now be forced to ignore errors** caused by invalid value formats in some columns, courtesy of the new parameter `on_error`. And the parameter `log_verbosity` will display the skipped values as NOTICE-level messages. A future enhancement may allow pushing invalid values into a separate table.

- 16 **You can now alter generated table expressions** using the  
ii command ALTER COLUMN with the SET EXPRESSION clause.  
Previously, the only option was to remove it.

**The abbreviation AT LOCAL** for the current time zone can now be used instead of AT TIME ZONE with a time zone value, as defined in the SQL standard.

## Functions and types

The overloaded function random now accepts two parameters to define **the minimum and maximum values for the random number**. Supports int, bigint and numeric values.

The type interval now supports **infinite values**.

In addition to to\_hex for hexadecimal, similar functions are added **for binary (to\_bin) and octal (to\_oct) systems**.

**New format masks for the to\_timestamp function:** TZ (for Time Zone) and OF (Offset from UTC). Previously, you had to use to\_char.

The function xmltext, as defined in the SQL Standard, **converts the input string to XML**, properly escaping special characters.

**New Unicode functions:** unicode\_assigned checks if every character in a string has a valid Unicode value, unicode\_version returns the Unicode version for PostgreSQL, and icu\_unicode\_version returns the ICU version.

**Parameter names** have been assigned to some aggregate function definitions with more than one parameter.

**The glorious epic that is implementation of the 2016 SQL/JSON Standard has now concluded.** The missing functions `JSON_EXISTS`, `JSON_QUERY`, `JSON_VALUE` and `JSON_TABLE` are finally in, and the `jsonpath` language gets new methods for data conversion into `bigint`, `boolean`, `date`, `decimal`, `integer`, `number`, `string`, `time_tz`, `time`, `timestamp_tz` and `timestamp`.

## Logical replication

**Upgrading is much easier now.** The `pg_upgrade` tool migrates replication slots to the publisher, so subscribers just have to update the connection string. And when upgrading the subscriber, all subscriptions stay intact and replication continues without the need to resynchronize.

**A logical replica can use a physical one as a base:** it already has all the tables synchronized, so setting up the replication will not take long. The new tool `pg_createsubscriber` switches the physical replica into read-write mode, creates the publisher and subscriber pairs in one or several databases on both servers, and specifies the position to continue replicating from in the subscriptions' properties.

**Logical slots are migrated to replicas** and kept up to date. After switchover, all you have to do is update the connection string to point to the new server, and the replication resumes.

**You can check what replication workers are doing** in the `pg_stat_subscription` view's new column `worker_type`.

- 18 Possible values are: apply, parallel apply, and table synchronization.  
ii

The columns `conflicting` and `invalidation_reason` in the `pg_stat_replication` view show if and why **the logical replication slot became invalid**.

In addition to that, logical decoding and, consequently, logical replication have been **significantly optimized for cases with large numbers of subtransactions**.

## Vacuum

**Vacuum is now much less memory-intensive** thanks to the new dead tuples storage utilizing radix trees. Additionally, tuple ID lookup is faster, memory is allocated dynamically (used to be all `maintenance_work_mem` at once), and the 1 GB hard cap is no more. Now, most tables will be vacuumed in one pass, avoiding repeat index scans.

**Vacuum monitoring has improved.** There are two new columns in `pg_stat_progress_vacuum`: total number of indexes to be vacuumed (`indexes_total`) and the number of indexes already vacuumed (`indexes_processed`).

**WAL size is now smaller** in cases where rows are frozen during vacuuming. Both events are treated as one record in WAL.

## Backup and upgrade

**Incremental backup** on the page level has been added to PostgreSQL. The process `walsummarizer` (with the parameter `summarize_wal`) reads WAL and records which pages

are modified. The `pg_basebackup` tool's new parameter `--incremental` enables incremental backup and utilizes the new replication protocol features. `pg_combinebackup` is the recovery tool that assembles the full backup from the incremental backup and all the backups it depends on.

19  
ii

**You can specify the database to connect to** as the parameter `dbname` in `pg_basebackup` and `pg_receivewal`. The former (with the `-R` key) will also add the database name to `primary_conninfo` when writing to `postgresql.auto.conf`. Neither tool needs the database name as such, but it is required when connecting over a proxy and when setting up replication slot synchronization on a replica.

**New keys for `pg_dump`:** `--filter` specifies a file with a list of objects to be included into or excluded from the dump. `--exclude-extension` excludes a list of extensions.

**Quicker upgrades with `pg_upgrade`** for databases with multiple large objects and tables. `pg_dump` now groups the large objects in the backup index, and `pg_restore` commits changes in transactions of the size specified in `--transaction-size`. Additionally, for parallel recovery, the leader process logic is improved.

## Access control

**New privilege `MAINTAIN`** for `ANALYZE`, `VACUUM` (including `VACUUM FULL`), `CLUSTER`, `REINDEX`, `REFRESH MATERIALIZED VIEW` and `LOCK TABLE`. **Membership in `pg_maintain`** grants this privilege for every relation in the database.

- 20 **Direct connection over TLS** is now available for client applications using libpq. Enabled by setting the parameter `sslnegotiation` to `requiredirect`.
- ii

**New default privilege display:** `psql` returns `(none)` if there are no privileges, and also respects `\pset null`.

## Query planning and execution

**Materialized CTE statistics** are now available to the planner. This includes CTE column statistics and row sorting order.

The planner can **swap tables in GROUP BY**, for example, to utilize index or incremental sorting.

**Functions with subtransactions are now PARALLEL SAFE.** This includes, in particular, PL/pgSQL functions with `EXCEPTION` blocks. Additionally, the parallel section of the plan can also include **InitPlan nodes**.

Improvements for the `EXPLAIN` command:

- The parameter `memory` displays **how much memory** was used to design the plan.
- The parameter `serialize` shows the **cost of conversion** of the query result into a text or binary format to be sent to the client (particularly for `TOAST` assembly).
- Improved readability of **SubPlan and InitPlan nodes**.

Many indexing features have always been available for B-trees only. Now, other index types can take advantage of them too:

- Hash indexes allows for **identifying altered rows** on the logical replication subscriber.

- **Incremental sorting** works with GiST and SP-GiST indexes. 21
- BRIN indexes can be created **in parallel**. ii

Some partitioning improvements:

- You can now set **exclusion constraints**.
- Full support for **GENERATED AS IDENTITY columns**.

## Monitoring

The new view `pg_stat_checkpoint` displays **checkpoint process statistics**, previously available in `pg_stat_bgwriter`. It also shows **replica restart points statistics**. Columns `buffers_backend` and `buffers_backend_fsync` are gone from `pg_stat_bgwriter`, since `pg_stat_io` now shows more accurate data about the processes.

`pg_stat_statements` view:

- New columns tracking **full statistics reset times** and **minimum and maximum time statistics** (`stats_since` and `minmax_stats_since`).
- **Normalization of commands** CALL, PREPARE TRANSACTION, COMMIT/ROLLBACK PREPARED and DEALLOCATE: their parameters are replaced with constant values, so the commands will be recorded as a single event.

**Range type column statistics** are now available in `pg_stats`.

Extension developers can now **define their own wait events**. The pioneers are the people behind the `dblink` and `postgres_fdw` extensions.

**Wait event types and descriptions** are available in the new view `pg_wait_events` as well as in the documentation.

## Event triggers

You can finally create an ON LOGIN trigger **for database connection events**.

Event triggers **work for the REINDEX command**.

The parameter *event\_triggers* **enables or disables event triggers**, primarily for debugging purposes.

## Configuration parameters

The parameter *old\_snapshot\_threshold*, first introduced in PostgreSQL 9.6, is gone, after issues with implementation and efficiency came up.

The new parameter *transaction\_timeout* **limits total transaction time**. Old *statement\_timeout* and *idle\_in\_transaction\_session\_timeout* parameters did not guarantee that a transaction will complete or terminate within a specific timeframe.

The parameter *huge\_pages\_status* shows if any **huge pages were allocated** if *huge\_pages* = try.

Some new configuration parameters define **SLRU cache sizes**, allowing for more fine-tuning:

- *commit\_timestamp\_buffers* – time of commit,
- *multixact\_member\_buffers* and *multixact\_offset\_buffers* – multitransactions,
- *subtransaction\_buffers* – nested transactions,
- *notify\_buffers* – asynchronous messages,

- *transaction\_buffers* – transaction statuses,
- *serializable\_buffers* – serializable transaction conflicts.

23  
ii

The ALTER SYSTEM command now can **record custom parameters** to postgresql.auto.conf. The parameter *allow\_alter\_system* can block ALTER SYSTEM from **accidentally altering the configuration**.

## Localization

New **stock provider** with support for C and C.UTF8 locales.

## Miscellaneous

You can use the key `--synch-method` to set the **disc synchronization method** for some utilities: `initdb`, `pg_basebackup`, `pg_checksums`, `pg_rewind`, `pg_upgrade` and `pg_dump`. The default method is `fsync`.

The `walsender` process **reads WAL records** from buffers when possible, bypassing the file system.

Previously, PostgreSQL sent a separate system call for each page and hoped that the OS would combine these operations on its end. Now, the **multiblock read infrastructure** exists: a sequence of pages (up to *io\_combine\_limit* bytes at a time) can be retrieved in one pass. Asynchronous reading is to be added in the future, enabling page processing as they come in, without waiting for the operation to complete.



# III Installation and Quick Start

What is required to get started with PostgreSQL? In this chapter, we'll learn how to install PostgreSQL and manage the corresponding service. In the next chapter, we'll continue by creating a simple database and trying out some basic SQL queries.

We are going to use a regular (often called “vanilla”) distribution of PostgreSQL 17. Depending on your operating system, the process of installing and setting up PostgreSQL will differ:

- If you are using Windows, read on.
- To set up PostgreSQL on Linux-based Debian or Ubuntu systems, go to p. 30.

For other operating systems, you can view installation instructions online: [www.postgresql.org/download](http://www.postgresql.org/download).

You can also use Postgres Pro Standard 17: it is fully compatible with vanilla PostgreSQL, includes some additional features developed by Postgres Professional, and is free when used for trial or educational purposes. Check out installation instructions at [postgrespro.com/products/download](http://postgrespro.com/products/download) in this case.

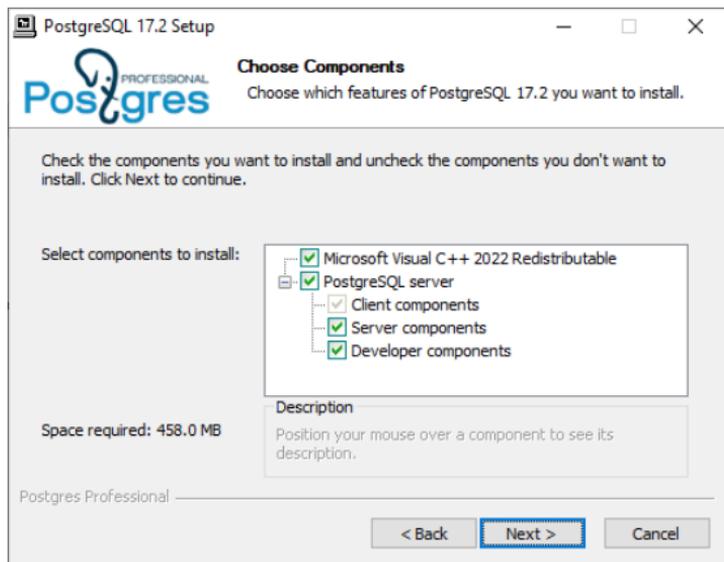
# Windows

## Installation

Download the PostgreSQL installer, launch it, and select the installation language: [postgrespro.com/windows](https://postgrespro.com/windows).

The installer provides a conventional wizard interface: you can simply keep clicking the “Next” button if you are fine with the default options. Let’s go over the main steps.

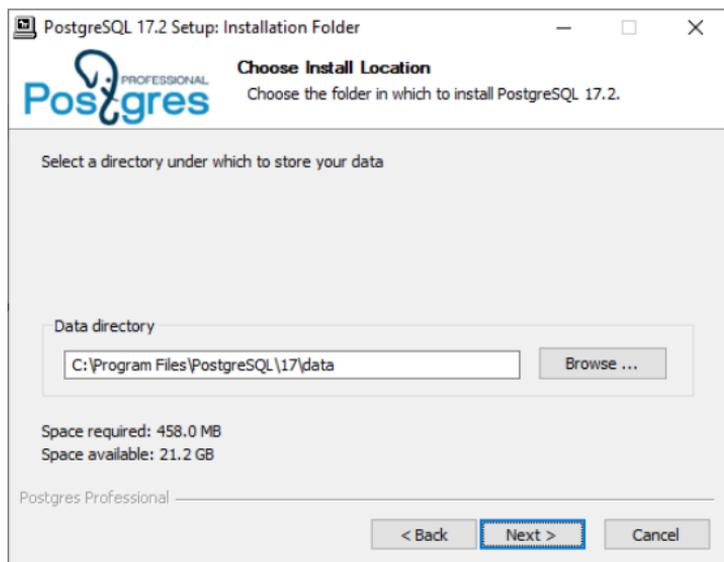
Choose components (keep the current selection if you are unsure of what to choose):



Then you have to specify PostgreSQL installation directory. By default, the PostgreSQL server is installed in `C:\Program Files\PostgreSQL\17`.

You can also specify the location of the data directory.

27  
iii

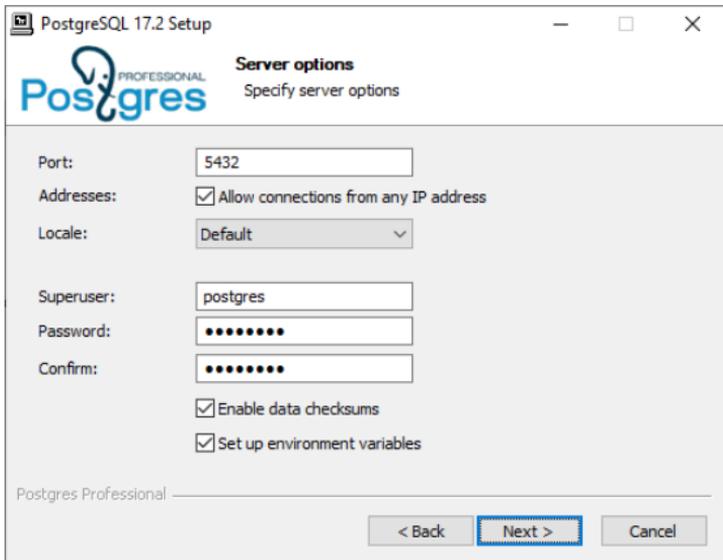


This directory will hold all the information stored in your database system, so make sure you have enough disk space if you plan to keep a lot of data.

If you plan to store data in a non-English language, ensure that you select the corresponding locale (or leave it as “Default” if your Windows locale settings are correct).

Enter and confirm the password for the postgres database user. You should also select the “Set up environment variables” checkbox to connect to the PostgreSQL server as the current OS user.

You can leave the default settings in all the other fields.

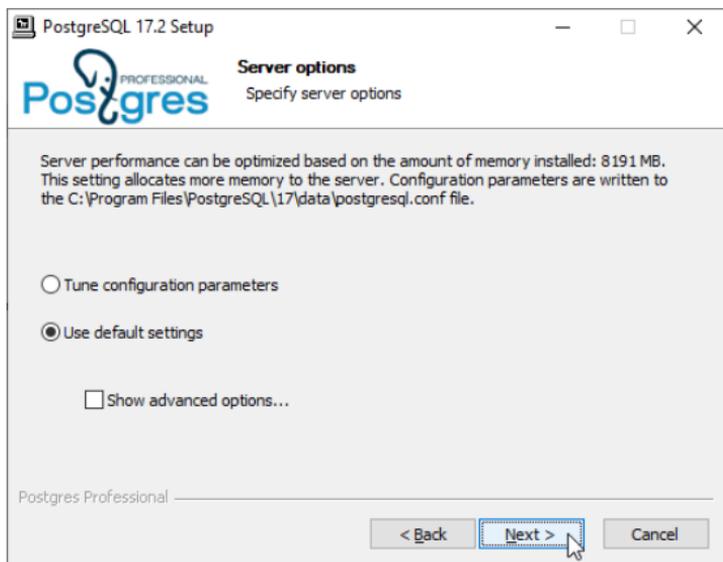


If you are planning to install PostgreSQL for training purposes only, you can select the “Use the default settings” option for the database system to take up less RAM.

## Managing the Service and the Main Files

When PostgreSQL is installed, the “postgresql-17” service is registered on your system. This service is launched automatically at the system startup under the Network Service account. If required, you can change the service settings using the standard Windows options.

To temporarily stop the service, run the “Stop Server” program from the Start menu subfolder that you have selected at installation time.



To start the service, run the “Start Server” program from the same folder.

If an error occurs when starting the service, you can view the server log to find out its cause. The log file is located in the log subdirectory of the database directory chosen at installation time (you can typically find it at C:\Program Files\PostgreSQL\17\data\log). Logging is regularly switched to a new file. You can find the required file either by the last modified date or by the filename that includes the date and time of the switchover to this file.

There are several important configuration files that define server settings. They are located in the database directory. You do not have to modify them to get started with PostgreSQL, but you’ll definitely need them in real work. Take a look into these files; they are fully documented:

- 30  
iii
- `postgresql.conf` is the main configuration file that contains server parameters.
  - `pg_hba.conf` defines access rules. For security reasons, the default configuration only allows local system access, requiring password authentication.

Now we are ready to connect to the database and try out some commands and SQL queries. Go to the “Trying SQL” chapter on p. 35.

## Debian and Ubuntu

### Installation

If you are using Linux, you need to add PGDG (PostgreSQL Global Development Group) package repository. At the moment, the supported Debian versions are 10 “Buster,” 11 “Bullseye,” and 12 “Bookworm.” The currently supported Ubuntu versions are 20.04 “Focal,” 22.04 “Jammy,” 23.10 “Mantic” and 24.04 “Noble”.

Run the following commands in the console window:

```
$ sudo apt-get install lsb-release
$ sudo sh -c 'echo "deb \
http://apt.postgresql.org/pub/repos/apt/ \
$(lsb_release-cs)-pgdg main" \
> /etc/apt/sources.list.d/pgdg.list'
$ wget --quiet -O - \
https://postgresql.org/media/keys/ACCC4CF8.asc \
| sudo apt-key add -
```

Once the repository is added, let's update the list of packages:

```
$ sudo apt-get update
```

Before starting the installation, check localization settings:

```
$ locale
```

If you plan to process non-English data, the `LC_CTYPE` and `LC_COLLATE` variables may have to be configured. For example, it makes sense to set these variables to “`fr_FR.UTF8`” for the French language, even though the “`en_US.UTF8`” may do too:

```
$ export LC_CTYPE=fr_FR.UTF8
```

```
$ export LC_COLLATE=fr_FR.UTF8
```

You should also make sure that the operating system has the required locale installed:

```
$ locale -a | grep fr_FR  
fr_FR.utf8
```

If this is not the case, generate the locale, as follows:

```
$ sudo locale-gen fr_FR.utf8
```

Now we can start the installation:

```
$ sudo apt-get install postgresql-17
```

It was the final step; once the installation command completes, PostgreSQL will be installed and launched. To check that the server is ready to use, run:

```
$ sudo -u postgres psql -c 'select version()'
```

If all went well, the current PostgreSQL version is returned.

## Managing the Service and the Main Files

When PostgreSQL is installed, a special `postgres` user is created on your system. All the server processes work on behalf of this user, and all the database files belong to this user as well. PostgreSQL will be started automatically at the operating system boot. It's not a problem with the default settings: if you are not working with the database server, it consumes very little of system resources. If you decide to disable automatic startup, run:

```
$ sudo systemctl disable postgresql
```

To temporarily stop the database server service, enter:

```
$ sudo systemctl stop postgresql
```

You can launch the server service as follows:

```
$ sudo systemctl start postgresql
```

You can also check the current state of the service:

```
$ sudo systemctl status postgresql
```

If the service fails to start, check the server log for details. Take a closer look at the latest log entries in `/var/log/postgresql/postgresql-17-main.log`.

All information stored in the database is located in the `/var/lib/postgresql/17/main/` directory. If you are going to keep a lot of data, make sure that you have enough disk space.

Server settings are defined by several configuration files. There's no need to edit all these files to get started, but it's worth checking them out since you'll definitely need them in the future:

33  
iii

- `/etc/postgresql/17/main/postgresql.conf` is the main configuration file that contains server parameters.
- `/etc/postgresql/17/main/pg_hba.conf` file defines access settings. For security reasons, the default configuration only allows access from the local system on behalf of the database user that has the same name as the current OS user.

Now it's time to connect to the database and try out SQL.



# IV Trying SQL

## Connecting via psql

To connect to the database server and start executing commands, you need to have a client application. In the “PostgreSQL for Applications” chapter, we will talk about sending queries from applications written in different programming languages. And here we’ll explain how to work with the psql client from the command line in interactive mode.

Unfortunately, many people are not very fond of the command line nowadays. Yet it is really worth mastering.

First of all, psql is a standard client application included in all PostgreSQL packages, so it’s always available. Having a customized environment is beneficial, but there is no need to get lost on an unfamiliar system.

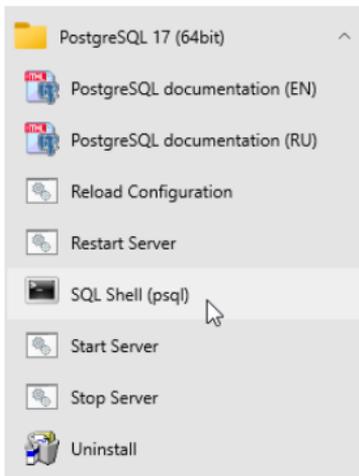
Secondly, psql is really convenient for everyday DBA tasks, writing small queries, and automating processes. For example, you can use it to periodically deploy application code updates on your database server. The psql client provides its own commands that can help you find your way around database objects and display the data stored in tables in a convenient format.

However, if you are used to working in graphical user interfaces, try pgAdmin (we’ll get back to it later) or other simi-

lar products: [wiki.postgresql.org/wiki/Community\\_Guide\\_to\\_PostgreSQL\\_GUI\\_Tools](http://wiki.postgresql.org/wiki/Community_Guide_to_PostgreSQL_GUI_Tools).

To start `psql` on a Linux system, run this command:

```
$ sudo -u postgres psql
```



On Windows, open the Start menu and launch the “SQL Shell (`psql`)” program. When prompted, enter the password for the `postgres` user that you set when installing PostgreSQL.

Windows users may run into encoding issues when viewing non-Latin characters in the terminal. If that is the case, make sure that a TrueType font is selected in the properties of the terminal window (typically, “Lucida Console” or “Consolas”).

As a result, you should see the same prompt on both operating systems: `postgres=#`. Here “`postgres`” is the name of the database to which you are connected right now. A single PostgreSQL server can serve several databases, but you can work only with one of them at a time.

Now let’s try out some commands. Enter only the part printed in bold; the prompt and the system response are provided here solely for your convenience.

Let's create a new database called test:

```
postgres=# CREATE DATABASE test;  
CREATE DATABASE
```

Don't forget to finish each command with a semicolon: PostgreSQL expects you to continue typing until you enter this symbol (so you can split the command between multiple lines).

Now let's connect to the created database:

```
postgres=# \c test  
You are now connected to database "test" as user  
"postgres".  
test=#
```

As you can see, the command prompt has changed to test=#.

The command that we've just entered does not look like SQL, as it starts with a backslash. This is a convention for special commands that can only be run in psql. If you are using pgAdmin or another GUI tool, skip all commands starting with a backslash or find an equivalent.

There are quite a few psql commands, and we'll use some of them a bit later. To get the full list of psql commands right now, you can run:

```
test=# \?
```

Since the reference information is quite bulky, it will be displayed in a pager program of your operating system, which is usually more or less.

## Tables

Relational database management systems present data as **tables**. The structure of the table is defined by its **columns**; the data itself is stored in table **rows**. The data is not ordered, so rows are not necessarily stored in the same order they were added to the table.

For each column, a **data type** is defined. All the values in the corresponding row fields must belong to this type. You can use multiple built-in data types provided by PostgreSQL ([postgrespro.com/doc/datatype](https://www.postgresql.org/docs/datatype)) or add your own custom types, but here we'll cover just a few main ones:

- integer
- text
- boolean, which is a logical data type taking true or false values

Apart from regular values defined by the data type, a field can be NULL. It can be interpreted as “the value is unknown” or “the value is not set.”

Let's create a table of university courses:

```
test=# CREATE TABLE courses(  
test(#  c_no text PRIMARY KEY,  
test(#  title text,  
test(#  hours integer  
test(# );  
CREATE TABLE
```

Note that the `psql` command prompt has changed: it is a hint that the command continues on the new line. For convenience, we will not repeat the prompt on each line in the examples that follow.

The above command creates the `courses` table with three columns: `c_no` specifies the course number represented as a text string, `title` provides the course title, and `hours` lists an integer number of lecture hours.

Apart from columns and data types, we can define integrity constraints that will be checked automatically: PostgreSQL will not allow invalid data in the database. In this example, we have added the PRIMARY KEY constraint for the `c_no` column. It means that all the values in this column must be unique, and NULLs are not allowed. Such a column can be used to distinguish one table row from another. The [postgrespro.com/doc/ddl-constraints](https://www.postgresql.com/doc/ddl-constraints) page lists all the available constraints.

You can find the exact syntax of the CREATE TABLE command in the documentation, or view command-line help right in `psql`:

```
test=# \help CREATE TABLE
```

Such reference information is available for each SQL command. To get the full list of SQL commands, run `\help` without arguments.

## Filling Tables with Data

Let's insert some rows into the created table:

```
test=# INSERT INTO courses(c_no, title, hours)
VALUES ('CS301', 'Databases', 30),
       ('CS305', 'Networks', 60);
```

```
INSERT 0 2
```

40 If you need to perform a bulk data upload from an external  
iv source, the INSERT command is not the best choice. Instead,  
you can use the COPY command specifically designed for this  
purpose: [postgrespro.com/doc/sql-copy](http://postgrespro.com/doc/sql-copy).

We'll need two more tables for subsequent examples: students and exams.

For each student, we are going to store their name and the year of admission (start year). The student ID number will serve as the student's identifier.

```
test=# CREATE TABLE students(  
      s_id integer PRIMARY KEY,  
      name text,  
      start_year integer  
);
```

```
CREATE TABLE
```

```
test=# INSERT INTO students(s_id, name, start_year)  
VALUES (1451, 'Anna', 2014),  
       (1432, 'Victor', 2014),  
       (1556, 'Nina', 2015);
```

```
INSERT 0 3
```

The exams table contains all the scores received by students in the corresponding course. Thus, students and courses are connected by the many-to-many relationship: each student can take exams in multiple courses, and each exam can be taken by multiple students.

Each table row is uniquely identified by the combination of a student ID and a course number. Such an integrity constraint pertaining to several columns at once is defined by the CONSTRAINT clause:

```
test=# CREATE TABLE exams(  
    s_id integer REFERENCES students(s_id),  
    c_no text REFERENCES courses(c_no),  
    score integer,  
    CONSTRAINT pk PRIMARY KEY(s_id, c_no)  
);  
CREATE TABLE
```

Besides, the REFERENCES clause here defines two referential integrity checks called **foreign keys**. Such keys show that the values of one table **reference** rows of another table.

After any action performed on the database, PostgreSQL checks that all the s\_id identifiers in the exams table correspond to real students (that is, entries in the students table), while c\_no course numbers correspond to real courses. Thus, it is impossible to assign a score on a non-existing subject or to a non-existent student, regardless of the user actions or possible application errors.

Let's assign several scores to our students:

```
test=# INSERT INTO exams(s_id, c_no, score)  
VALUES (1451, 'CS301', 5),  
       (1556, 'CS301', 5),  
       (1451, 'CS305', 5),  
       (1432, 'CS305', 4);  
INSERT 0 4
```

## Data Retrieval

### Simple Queries

To read data from tables, use the SQL operator SELECT. For example, let's display two columns of the courses table. The AS clause allows you to rename the column if required:

```
42 test=# SELECT title AS course_title, hours
iv FROM courses;
   course_title | hours
-----+-----
   Databases   |    30
   Networks    |    60
(2 rows)
```

The asterisk \* displays all the columns:

```
test=# SELECT * FROM courses;
  c_no | title      | hours
-----+-----+-----
  CS301 | Databases |    30
  CS305 | Networks  |    60
(2 rows)
```

In production applications, it is recommended to explicitly specify only those columns that are really needed: then the query is performed more efficiently, and the result does not depend on new columns that may appear. But in interactive queries you can simply use an asterisk.

The result can contain several rows with the same data. Even if all rows in the original table are different, the data can appear duplicated if not all the columns are displayed:

```
test=# SELECT start_year FROM students;
 start_year
-----
        2014
        2014
        2015
(3 rows)
```

To select all **different** start years, specify the **DISTINCT** keyword after **SELECT**:

```
test=# SELECT DISTINCT start_year FROM students;
 start_year
-----
         2014
         2015
(2 rows)
```

For details, see the documentation: [postgrespro.com/doc/sql-select#SQL-DISTINCT](http://postgrespro.com/doc/sql-select#SQL-DISTINCT).

In general, you can use any expressions after the SELECT operator. And if you omit the FROM clause, the query will return a single row. For example:

```
test=# SELECT 2+2 AS result;
 result
-----
         4
(1 row)
```

When you select some data from a table, it is usually required to return only those rows that satisfy a certain condition. This filtering condition is specified in the WHERE clause:

```
test=# SELECT * FROM courses WHERE hours > 45;
 c_no | title | hours
-----+-----+-----
 CS305 | Networks |    60
(1 row)
```

The condition must be of a logical type. For example, it can contain operators =, <> (or !=), >, >=, <, <=, as well as combine simple conditions using logical operations AND, OR, NOT, and parenthesis (like in regular programming languages).

Handling NULLs is a bit more subtle. The result will contain only those rows for which the filtering condition is true; if the condition is false or **undefined**, the row is excluded.

44 Remember:

iv

- The result of comparing something to NULL is undefined.
- The result of logical operations on NULLs is usually undefined (exceptions: true OR NULL = true, false AND NULL = false).
- To check whether the value is undefined, the following operators are used: IS NULL (IS NOT NULL) and IS DISTINCT FROM (IS NOT DISTINCT FROM).

The coalesce expression is often used to replace NULL values with something else, such as an empty string for text types or zero for numeric types.

For more details, see the documentation: [postgrespro.com/doc/functions-comparison](https://www.postgresql.org/docs/functions-comparison).

## Joins

A well-designed database should not contain redundant data. For example, the exams table must not contain student names, as this information can be found in another table by the number of the student ID card.

For this reason, to get all the required values in a query, it is often necessary to join the data of several tables, specifying their names in the FROM clause:

```
test=# SELECT * FROM courses, exams;
```

c_no	title	hours	s_id	c_no	score
CS301	Databases	30	1451	CS301	5
CS305	Networks	60	1451	CS301	5
CS301	Databases	30	1556	CS301	5

CS305	Networks		60	1556	CS301		5
CS301	Databases		30	1451	CS305		5
CS305	Networks		60	1451	CS305		5
CS301	Databases		30	1432	CS305		4
CS305	Networks		60	1432	CS305		4

(8 rows)

This result is called the direct or Cartesian product of tables: each row of one table is appended to each row of the other table.

As a rule, you can get a more useful and informative result if you specify the join condition in the WHERE clause. Let's match the courses to the corresponding exams to get all the scores for all the courses:

```
test=# SELECT courses.title, exams.s_id, exams.score
FROM courses, exams
WHERE courses.c_no = exams.c_no;
```

title	s_id	score
Databases	1451	5
Databases	1556	5
Networks	1451	5
Networks	1432	4

(4 rows)

Another way to join tables is to explicitly use the JOIN keyword. Let's display all the students and their scores for the "Networks" course:

```
test=# SELECT students.name, exams.score
FROM students
JOIN exams
ON students.s_id = exams.s_id
AND exams.c_no = 'CS305';
```

46  
iv

name	score
Anna	5
Victor	4

(2 rows)

From the database point of view, these queries are completely equivalent, so you can use any approach that seems more natural.

In this example, the result does not include any rows of the table specified on the left side of the join clause if they have no pair in the right table: although the condition is applied to the subjects, the students that did not take the exam in this subject are also excluded. To include all the students into the result, we have to use the outer join:

```
test=# SELECT students.name, exams.score
FROM students
LEFT JOIN exams
  ON students.s_id = exams.s_id
  AND exams.c_no = 'CS305';
```

name	score
Anna	5
Victor	4
Nina	

(3 rows)

Note that the rows of the left table that don't have a counterpart in the right table are added to the result (that's why the operation is called LEFT JOIN). The corresponding values of the right table are NULL in this case.

The WHERE conditions are applied to the result of the join operation. Thus, if you move the subject restriction from the join condition to the WHERE clause, Nina will be excluded from the result because the corresponding exams.c\_no is NULL:

```
test=# SELECT students.name, exams.score
FROM students
LEFT JOIN exams ON students.s_id = exams.s_id
WHERE exams.c_no = 'CS305';
```

```
   name | score
-----+-----
 Anna   |     5
 Victor |     4
(2 rows)
```

Don't be afraid of joins. It is a common operation for database management systems, and PostgreSQL has a whole range of efficient mechanisms to perform it. Do not join data at the application level; let the database server handle it properly.

For more details, see the documentation: [postgrespro.com/doc/sql-select#SQL-FROM](https://www.postgresql.org/docs/sql-select.html).

## Subqueries

The SELECT operation returns a table, which can be displayed as the query result (as we have already seen) or used in another SQL query. Such a nested SELECT command in parentheses is called a **subquery**.

If a subquery returns exactly one row and one column, you can use it as a regular scalar expression:

```
test=# SELECT name,
      (SELECT score
       FROM exams
       WHERE exams.s_id = students.s_id
       AND exams.c_no = 'CS305')
FROM students;
```

48  
iv

name	score
Anna	5
Victor	4
Nina	

(3 rows)

If a scalar subquery used in the list of SELECT expressions does not contain any rows, NULL is returned (as in the last row of the result in the example above). Thus, you can expand scalar subqueries by replacing them with a join, but it must be an outer join.

Scalar subqueries can also be used in filtering conditions. Let's display all the exams taken by the students enrolled after 2014:

```
test=# SELECT *
FROM exams
WHERE (SELECT start_year FROM students
       WHERE students.s_id = exams.s_id) > 2014;
```

s_id	c_no	score
1556	CS301	5

(1 row)

You can also add filtering conditions to subqueries returning an arbitrary number of rows. SQL offers several predicates for this purpose. For example, IN checks whether the table returned by the subquery contains the specified value.

Let's display all the students who have any scores in the specified course:

```
test=# SELECT name, start_year
FROM students
WHERE s_id IN (SELECT s_id FROM exams
              WHERE c_no = 'CS305');
```

name	start_year
Anna	2014
Victor	2014

(2 rows)

There is also the NOT IN form of this predicate, which returns the opposite result. For example, the following query returns the list of students who did not get any excellent scores:

```
test=# SELECT name, start_year
FROM students
WHERE s_id NOT IN (SELECT s_id
                   FROM exams
                   WHERE score = 5);
```

name	start_year
Victor	2014

(1 rows)

Note that this query result can also include those students who have not received any scores at all.

Another option is to use the EXISTS predicate, which checks whether the subquery returns at least one row. With this predicate, you can rewrite the previous query as follows:

```
test=# SELECT name, start_year
FROM students
WHERE NOT EXISTS (SELECT s_id
                  FROM exams
                  WHERE exams.s_id = students.s_id
                  AND score = 5);
```

name	start_year
Victor	2014

(1 rows)

50 For more details, see the documentation: [postgrespro.com/doc/functions-subquery](https://postgrespro.com/doc/functions-subquery).  
iv

In the examples above, we appended table names to column names to avoid ambiguity, but it is not always sufficient. For example, the same table can be used in the query twice, or the FROM clause can contain a nameless subquery instead of a table name. In such cases, you can specify an arbitrary name after the query, which is called an alias. Regular tables can also be assigned an alias.

Let's display student names and their scores for the "Databases" course:

```
test=# SELECT s.name, ce.score
FROM students s
JOIN (SELECT exams.*
      FROM courses, exams
      WHERE courses.c_no = exams.c_no
      AND courses.title = 'Databases') ce
ON s.s_id = ce.s_id;
```

name	score
Anna	5
Nina	5

(2 rows)

Here "s" is a table alias, while "ce" is a subquery alias. You should choose an alias that is short but comprehensive.

The same query can also be written without subqueries:

```
test=# SELECT s.name, e.score
FROM students s, courses c, exams e
WHERE c.c_no = e.c_no
AND c.title = 'Databases'
AND s.s_id = e.s_id;
```

## Sorting

As we already know, table data is not sorted. To return the rows in a particular order, we can use the `ORDER BY` clause with the list of sorting expressions. After each expression (sorting key), you can specify the sort direction: `ASC` for ascending (used by default), `DESC` for descending.

```
test=# SELECT * FROM exams
ORDER BY score, s_id, c_no DESC;
```

s_id	c_no	score
1432	CS305	4
1451	CS305	5
1451	CS301	5
1556	CS301	5

(4 rows)

Here the rows are first sorted by the score, in ascending order. For the same scores, the rows are further sorted by the student ID card number, also in ascending order. If the first two keys are the same, rows are additionally sorted by the course number, in descending order.

It makes sense to sort data at the end of the query, right before returning the result; this operation is usually unnecessary in subqueries.

For more details, see the documentation: [postgrespro.com/doc/sql-select#SQL-ORDERBY](https://www.postgresql.org/docs/sql-select#SQL-ORDERBY).

## Grouping

When grouping is used, the query returns summary data calculated from multiple rows stored in the original tables.

52 Grouping typically involves **aggregate functions**. For exam-  
iv ple, this is how we can display the total number of exams  
taken, the number of students who passed the exams, and  
the average score:

```
test=# SELECT count(*), count(DISTINCT s_id),  
avg(score)  
FROM exams;  
 count | count |      avg  
-----+-----+-----  
      4 |      3 | 4.7500000000000000  
(1 row)
```

You can get similar information by the course number if you  
provide the GROUP BY clause with grouping keys:

```
test=# SELECT c_no, count(*),  
count(DISTINCT s_id), avg(score)  
FROM exams  
GROUP BY c_no;  
 c_no | count | count |      avg  
-----+-----+-----+-----  
CS301 |      2 |      2 | 5.0000000000000000  
CS305 |      2 |      2 | 4.5000000000000000  
(2 rows)
```

For the full list of aggregate functions, see [postgrespro.com/doc/functions-aggregate](https://www.postgresql.org/docs/functions-aggregate).

In queries that use grouping, you may need to filter the rows  
based on the aggregation results. You can define such condi-  
tions in the HAVING clause. While the WHERE conditions are  
applied before grouping (and can use the columns of the origi-  
nal tables), the HAVING conditions take effect after grouping  
(so they can also use the columns of the resulting table).

Let's select the names of the students who got more than  
one excellent score (5), in any course:

```
test=# SELECT students.name
FROM students, exams
WHERE students.s_id = exams.s_id AND exams.score = 5
GROUP BY students.name
HAVING count(*) > 1;
 name
-----
 Anna
(1 row)
```

For more details, see the documentation: [postgrespro.ru/doc/sql-select#SQL-GROUPBY](http://postgrespro.ru/doc/sql-select#SQL-GROUPBY).

## Changing and Deleting Data

To modify data in a table, you should use the UPDATE operator, which provides new field values for rows defined by the WHERE clause (similar to the SELECT operator).

For example, let's double the number of lecture hours for the "Databases" course:

```
test=# UPDATE courses
SET hours = hours * 2
WHERE c_no = 'CS301';
UPDATE 1
```

For more details, see the documentation: [postgrespro.com/doc/sql-update](http://postgrespro.com/doc/sql-update).

Similarly, the DELETE operator deletes the rows defined by the WHERE clause:

```
test=# DELETE FROM exams WHERE score < 5;
DELETE 1
```

## Transactions

Let's extend our database schema a little bit and distribute our students between groups. Each group must have a monitor (a student of the same group responsible for the students' activities). For this purpose, let's create the groups table:

```
test=# CREATE TABLE groups(  
      g_no text PRIMARY KEY,  
      monitor integer NOT NULL REFERENCES students(s_id)  
);  
CREATE TABLE
```

Here we have applied the NOT NULL constraint, which forbids using undefined values.

Now we need another column in the students table: the group number. Luckily, we can add a new column into the already existing table:

```
test=# ALTER TABLE students  
ADD g_no text REFERENCES groups(g_no);  
ALTER TABLE
```

Using the psql command, you can always check which columns are defined in the table:

```
test=# \d students  
      Table "public.students"  
  Column | Type | Modifiers  
-----+-----+-----  
 s_id   | integer | not null  
 name   | text    |  
 start_year | integer |  
 g_no   | text    |  
 ...
```

You can also get the list of all the tables available in the database:

55  
iv

```
test=# \d
```

```
                List of relations
 Schema | Name   | Type  | Owner
-----+-----+-----+-----
 public | courses | table | postgres
 public | exams  | table | postgres
 public | groups | table | postgres
 public | students | table | postgres
(4 rows)
```

Now let's create a new group called "A-101," move all the students into this group, and make Anna its monitor.

Note the following subtle point. We cannot create a group without a monitor, but neither can a student become the monitor of the group unless they are already a member of this group—it would make our data logically incorrect and inconsistent. Taken separately, these two operations make no sense: they must be performed simultaneously. A **transaction** is a group of operations that form an indivisible logical unit of work.

So let's start our transaction:

```
test=# BEGIN;
```

```
BEGIN
```

Next, we need to add a new group, together with its monitor. Naturally, we cannot remember all the students' ID, so we'll use the following query right inside the command that adds new rows:

```
56 test=# INSERT INTO groups(g_no, monitor)
iv  SELECT 'A-101', s_id
    FROM students
    WHERE name = 'Anna';
INSERT 0 1
```

The asterisk in the prompt reminds us that the transaction is not yet completed.

Now let's open a new terminal window and launch another `psql` process: this session will be running in parallel with the first one. To avoid confusion, we will indent the commands of the second session.

Will the second session see the changes made in the first session?

```
postgres=# \c test
You are now connected to database "test" as user
"postgres".
test=# SELECT * FROM groups;
 g_no | monitor
-----+-----
(0 rows)
```

No, since the transaction is not yet completed.

To continue with our transaction, let's move all students to the newly created group:

```
test=# UPDATE students SET g_no = 'A-101';
UPDATE 3
```

The second session still gets consistent data, which was already present in the database when the uncommitted transaction was started.

```
test=# SELECT * FROM students;
 s_id | name  | start_year | g_no
-----+-----+-----+-----
 1451 | Anna  |         2014 |
 1432 | Victor |         2014 |
 1556 | Nina  |         2015 |
(3 rows)
```

Let's commit all our changes to complete the transaction:

```
test=# COMMIT;
COMMIT
```

Finally, the second session receives all the changes made by this transaction, as if they appeared all at once:

```
test=# SELECT * FROM groups;
 g_no | monitor
-----+-----
A-101 |    1451
(1 row)

test=# SELECT * FROM students;
 s_id | name  | start_year | g_no
-----+-----+-----+-----
 1451 | Anna  |         2014 | A-101
 1432 | Victor |         2014 | A-101
 1556 | Nina  |         2015 | A-101
(3 rows)
```

It is guaranteed that several important properties of the database system are always observed.

First of all, any transaction is executed either completely (like in the example above), or not at all. If at least one of the com-

58  
iv

mands results in an error, or we have aborted the transaction with the ROLLBACK command, the database stays in the same state as before the BEGIN command. This property is called **atomicity**.

Second, when a transaction is committed, all integrity constraints must hold true, otherwise the transaction has to be aborted. The data is consistent when the transaction starts, and it remains consistent at the end of the transaction, which gives this property its name: **consistency**.

Third, as the example has shown, other users will never see inconsistent data not yet committed by the transaction. This property is called **isolation**. Thanks to this property, the database system can serve multiple sessions in parallel, without sacrificing data consistency.

PostgreSQL is known for a very efficient implementation of isolation: several sessions can perform reads and writes in parallel, without blocking each other. Blocking occurs only if two different processes try changing the same row simultaneously.

And finally, **durability** is guaranteed: the committed data is never lost even in case of a failure (if the database is set up correctly and is regularly backed up, of course).

These properties are extremely important; it is impossible to imagine a relational database management system without them.

To learn more about transactions, see [postgrespro.com/doc/tutorial-transactions](https://www.postgresql.org/docs/tutorial-transactions) (Even more details are available here: [postgrespro.com/doc/mvcc](https://www.postgresql.org/docs/mvcc)).

## Useful psql Commands

59  
iv

<b>\?</b>	Command-line reference for psql.
<b>\h</b>	SQL Reference: the list of available commands or the syntax of a particular command.
<b>\x</b>	A switch that toggles between the regular table display (rows and columns) and an extended display (with each column printed on a separate line). This is useful for viewing several wide rows.
<b>\l</b>	List of databases.
<b>\du</b>	List of users.
<b>\dt</b>	List of tables.
<b>\di</b>	List of indexes.
<b>\dv</b>	List of views.
<b>\df</b>	List of functions.
<b>\dn</b>	List of schemas.
<b>\dx</b>	List of installed extensions.
<b>\dp</b>	List of privileges.
<b>\d name</b>	Detailed information about the specified object.
<b>\d+ name</b>	Extended detailed information about the specified object.
<b>\timing on</b>	Displays operator execution time.

## Conclusion

We have covered only a small portion of what you need to know about PostgreSQL, but we hope you have found it easy to get started. The SQL language enables you to construct queries of various complexity, while PostgreSQL provides an effective implementation and high-quality support of the standard. Try it yourself and experiment!

And one more important psql command. To end the session, enter:

```
test=# \q
```

# V Demo Database

## About the Demo Database

### Overview

To get to grips with more complex queries, we need to create a more substantial database (with not just three but eight tables) and fill it up with some reasonable data.

We have chosen airline flights as the subject area. Our database contains statistics on all flights operated by a hypothetical airline within a specific timeframe. These scenarios must be familiar to anyone who has ever traveled by plane, but we'll explain everything anyway.

The database schema is shown on p. 63. We aimed for a balance: keeping the database schema simple without unnecessary complexity, yet detailed enough to support interesting and meaningful queries.

The main entity in our schema is a **booking** (mapped to the bookings table). Each booking can include several passengers, with a separate **ticket** issued for each passenger (tickets). We do not have any reliable unique ID for a passenger as a person (who might have flown with our airline multiple times), so the passenger does not constitute a separate entity. We will assume that all the passengers are unique.

62 Each ticket always contains one or more **flight segments**  
v (ticket\_flights). A single ticket can include multiple segments if there are no direct flights between the departure and arrival airports or if it is a round-trip ticket.

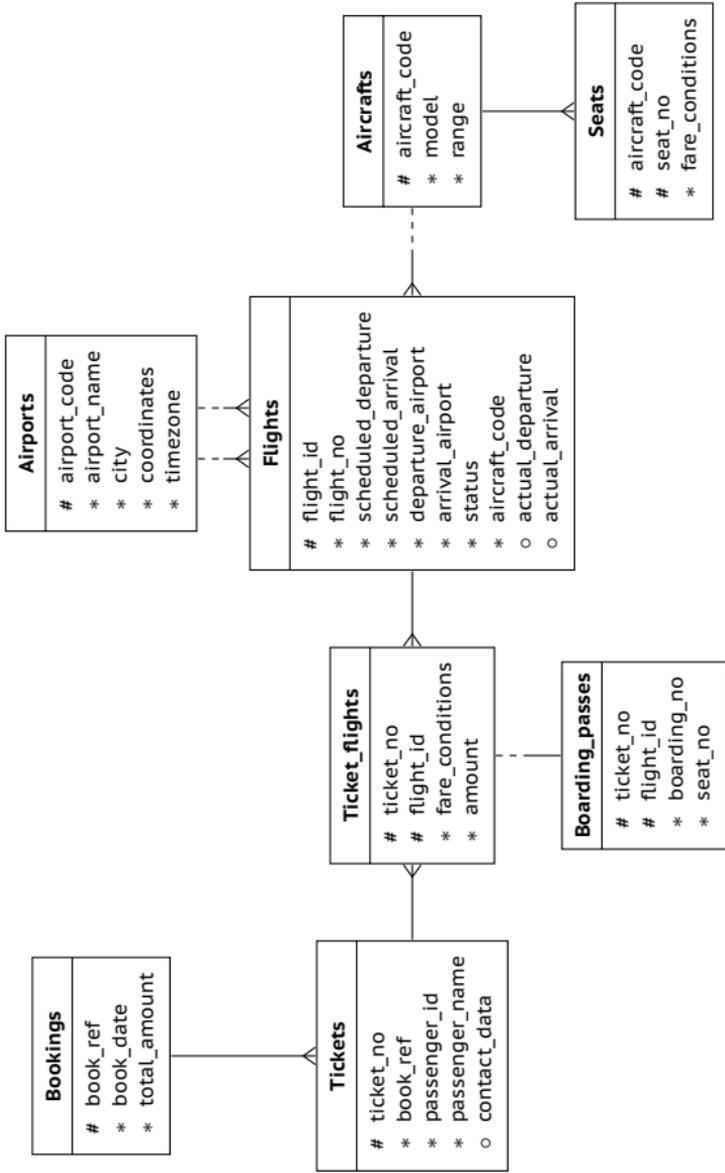
It is assumed that all tickets in a single booking have the same flight segments, even though there is no such constraint in the schema.

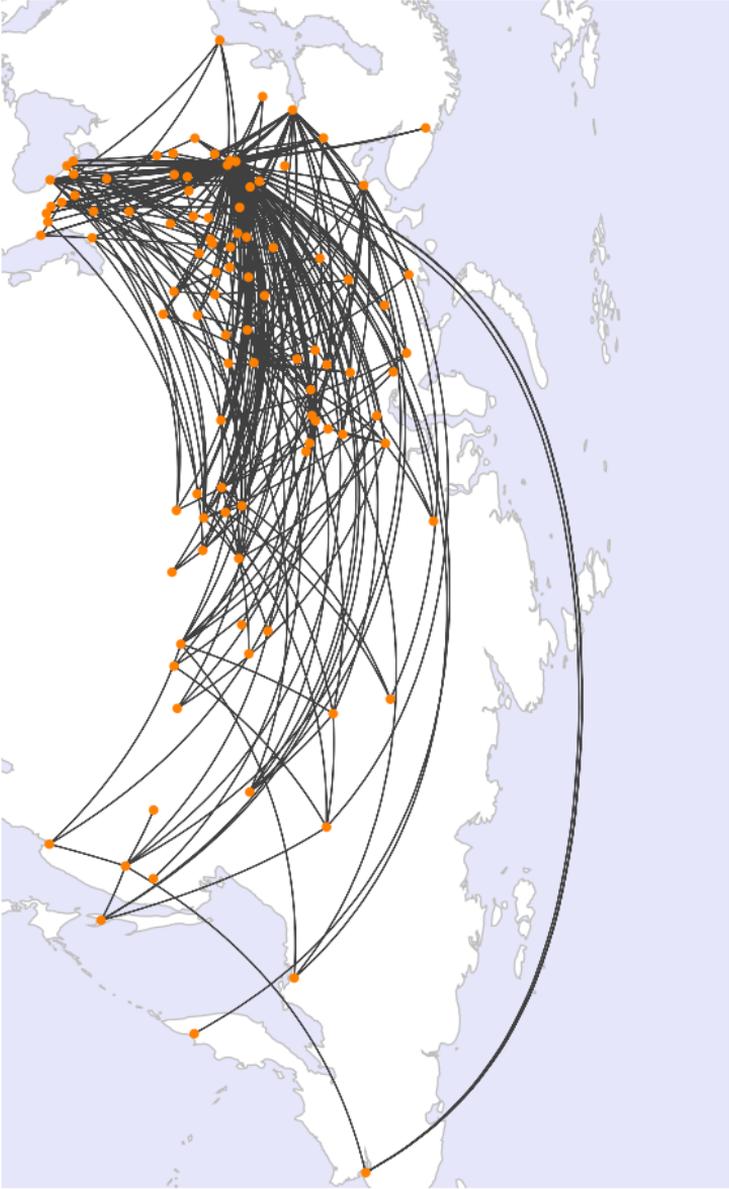
Each **flight** (flights) goes from one **airport** (airports) to another. Flights with the same flight number have the same points of departure and destination but different departure dates.

At flight check-in, each passenger is issued a **boarding pass** (boarding\_passes), where the seat number is specified. Passengers can check in for the flight only if they have a ticket for this flight. The flight/seat combination must be unique to avoid issuing several boarding passes for the same seat.

The number of **seats** in the aircraft and their distribution between different travel classes depend on the specific model of the **aircraft** performing the flight. It is assumed that each aircraft model has only one cabin configuration. The database schema does not include any checks on whether the seat specified in the boarding pass is actually available in the particular aircraft.

In the sections that follow, we'll describe each of the tables, as well as additional views and functions. You can always use the \d+ command to get the exact definition of any table, including data types and column descriptions.





## Bookings

Passengers must book their tickets in advance. The booking date (`book_date`) must be within one month of the flight. The booking is identified by its number (`book_ref`, a six-position combination of letters and digits).

The `total_amount` field stores the total price of all tickets included into the booking, for all passengers.

## Tickets

A ticket has a unique number (`ticket_no`), which consists of 13 digits.

The ticket contains the number of the passenger's identity document (`passenger_id`), as well as the passenger's first and last names (`passenger_name`) and contact information (`contact_data`).

Since a passenger's ID or name may change (e.g., due to a name change or passport renewal), uniquely identifying all tickets belonging to a single passenger is not always possible. For simplicity, let's assume that all the passengers are unique.

## Flight Segments

A flight segment links a ticket to a flight and is identified by both the ticket number and the flight number.

Each flight segment has its price (`amount`) and travel class (`fare_conditions`).

A unique ID can be either natural (if it is related to real-life objects) or surrogate (if it is generated by the system, typically as an increasing sequence of numbers).

The natural composite key of the `flights` table consists of the flight number (`flight_no`) and the date of the departure (`scheduled_departure`). To make foreign keys that refer to this table a bit shorter, a surrogate key `flight_id` is used as the primary key.

Each flight departs from one airport (`departure_airport`) and arrives at another (`arrival_airport`).

There is no such entity as a “connecting flight”: if there are no direct flights from one airport to another, the ticket simply includes all the required flight segments.

Each flight has a scheduled date and time of departure and arrival (`scheduled_departure` and `scheduled_arrival`). However, actual departure and arrival times (`actual_departure` and `actual_arrival`) may vary due to delays, which can range from a few minutes to several hours.

Flight status can take one of the following values:

- **Scheduled**  
The flight can be booked. This value is set one month before the planned departure date; at this point, the information about the flight is entered into the database.
- **On Time**  
The flight is open for check-in (twenty-four hours before the scheduled departure) and is not delayed.

- **Delayed**  
The flight is open for check-in (twenty-four hours before the scheduled departure), but is delayed.
- **Departed**  
The aircraft has already departed and is airborne.
- **Arrived**  
The aircraft has reached the point of destination.
- **Cancelled**  
The flight is cancelled.

## Airports

An airport is identified by a three-letter `airport_code` and has an `airport_name`.

There is no separate entity for the city; a city name is simply an airport attribute, which is required to identify all the airports of the same city. The table also includes coordinates (longitude and latitude) and the timezone.

## Boarding Passes

At the time of check-in, which opens twenty-four hours before the scheduled departure, the passenger is issued a boarding pass. Each boarding pass is uniquely identified by a combination of ticket and flight numbers.

Boarding pass numbers (`boarding_no`) are assigned sequentially based on check-in order. They are unique only within a specific flight. The boarding pass specifies the seat number (`seat_no`).

## Aircraft

Each aircraft model is identified by a three-character `aircraft_code`. The table also includes the name of the aircraft model and the maximum flying distance, in kilometers (range).

## Seats

Seats define the cabin configuration of each aircraft model. Each seat has a number (`seat_no`) and an assigned travel class (`fare_conditions`): Economy, Comfort, or Business.

## Flights View

There is a `flights_v` view built over the `flights` table. Views can be queried in the same way as tables, but they do not store any data: they simply perform a particular query. The following `psql` command displays the definition of the view and its query:

```
postgres=# \d+ flights_v
```

This view adds the following information:

- details about the airport of departure  
`departure_airport`, `departure_airport_name`,  
`departure_city`
- details about the airport of arrival  
`arrival_airport`, `arrival_airport_name`,  
`arrival_city`
- local departure time  
`scheduled_departure_local`, `actual_departure_local`

- local arrival time  
scheduled\_arrival\_local, actual\_arrival\_local
- flight duration  
scheduled\_duration, actual\_duration

## Routes View

The flights table contains some redundancies, which you can use to get route information that does not depend on the exact flight dates (flight number, airports of departure and destination, aircraft model).

This information constitutes the routes view. Besides, this view shows the days\_of\_week array representing days of the week on which flights are performed, and the planned flight duration.

## The “now” Function

The demo database contains a snapshot of data, similar to a backup of a real system captured at some point in time. For example, if a flight has the Departed status, it means that the aircraft was airborne at the time the backup was taken.

The snapshot time is saved in the bookings.now function. You can use this function in demo queries for cases that would typically require calling the now function.

Besides, the return value of this function determines the version of the demo database. The latest version, as of August 15, 2017, contains data from that date.

## Installation

### Installation from the Website

The demo database is available in three flavors, which differ only in the data size:

- [edu.postgrespro.com/demo-small-en.zip](https://edu.postgrespro.com/demo-small-en.zip)  
A small database with flight data for one month (21 MB, DB size is 280 MB).
- [edu.postgrespro.com/demo-medium-en.zip](https://edu.postgrespro.com/demo-medium-en.zip)  
A medium database with flight data for three months (62 MB, DB size is 702 MB).
- [edu.postgrespro.com/demo-big-en.zip](https://edu.postgrespro.com/demo-big-en.zip)  
A large database with flight data for one year (232 MB, DB size is 2638 MB).

A small database is quite suitable for learning how to write queries, but if you would like to deal with query optimization specifics, choose the large database: then you'll be able to see how queries work on large volumes of data.

The files contain a logical backup of the demo database created with the `pg_dump` utility. Note that if you already have some database named `demo`, it will be dropped and restored from this backup. The user who runs the script becomes the owner of this database.

To install the demo database on a Linux system, switch to the `postgres` user and download the corresponding file. For example, to download the small database, do the following:

```
$ sudo su - postgres
```

```
$ wget https://edu.postgrespro.com/demo-small-en.zip
```

Then run the following command:

```
$ zcat demo-small-en.zip | psql
```

On Windows, download the [edu.postgrespro.com/demo-small-en.zip](http://edu.postgrespro.com/demo-small-en.zip) file, double-click it to open the archive, and copy the `demo-small-en-20170815.sql` file into the `C:\Program Files\PostgreSQL\17` directory.

The pgAdmin application (described on p. 111) does not support restoring databases from this type of backup. So you should start psql (by clicking the “SQL Shell (psql)” shortcut) and run the following command:

```
postgres# \i demo-small-en-20170815.sql
```

If the file is not found, check the “Start in” property of the shortcut; the file must be located in this directory.

## Sample Queries

### A Couple of Words about the Schema

Once installation is complete, launch psql and connect to the demo database:

```
postgres=# \c demo
```

You are now connected to database "demo" as user "postgres".

The bookings schema contains all necessary entities. When connected to the database, this schema is used automatically, so there is no need to specify it explicitly:

72 demo=# **SELECT \* FROM aircrafts;**

v

aircraft_code	model	range
773	Boeing 777-300	11100
763	Boeing 767-300	7900
SU9	Sukhoi Superjet-100	3000
320	Airbus A320-200	5700
321	Airbus A321-200	5600
319	Airbus A319-100	6700
733	Boeing 737-300	4200
CN1	Cessna 208 Caravan	1200
CR2	Bombardier CRJ-200	2700

(9 rows)

However, we must specify the schema for the `bookings.now` function, as we have to differentiate it from the standard `now` function:

demo=# **SELECT bookings.now();**

```
          now
-----
2017-08-15 18:00:00+03
(1 row)
```

The following query returns cities and airports:

demo=# **SELECT airport\_code, city  
FROM airports LIMIT 4;**

airport_code	city
YKS	Yakutsk
MJZ	Mirnyj
KHV	Khabarovsk
PKC	Petropavlovsk

(4 rows)

The contents of the database is provided in English and in Russian. You can switch between these languages by setting the `bookings.lang` parameter to `en` or `ru`, respectively. By default, the English language is selected. You can change this setting as follows:

```
demo=# ALTER DATABASE demo SET bookings.lang = ru;
```

```
ALTER DATABASE
```

The language has been changed at the database level; now we have to reconnect to the database.

```
demo=# \c
```

You are now connected to database "demo" as user "postgres".

```
demo=# SELECT airport_code, city
FROM airports LIMIT 4;
```

airport_code	city
YKS	Якутск
MJZ	Мирный
KHV	Хабаровск
PKC	Петропавловск-Камчатский

(4 rows)

To understand how it works, take a look at the `aircrafts` or `airports` definition using the `\d+ psql` command.

If you want to learn more about managing schemas, see the documentation: [postgrespro.com/doc/ddl-schemas](https://postgrespro.com/doc/ddl-schemas).

For more details on setting configuration parameters, see [postgrespro.com/doc/config-setting](https://postgrespro.com/doc/config-setting).

## Simple Queries

Let's use this schema to discuss several problems, starting from the simplest questions and getting to more complex ones. Most problems are followed by a solution, but if you want to learn SQL effectively, try to write your own query before looking at the answer.

**Problem.** Who traveled from Moscow (SVO) to Novosibirsk (OVB) on seat 1A the day before yesterday, and when was the ticket booked?

**Solution.** "The day before yesterday" is counted from the `booking.now` value, not from the current date.

```
SELECT t.passenger_name,  
       b.book_date  
FROM   bookings b  
       JOIN tickets t  
         ON t.book_ref = b.book_ref  
       JOIN boarding_passes bp  
         ON bp.ticket_no = t.ticket_no  
       JOIN flights f  
         ON f.flight_id = bp.flight_id  
WHERE  f.departure_airport = 'SVO'  
AND    f.arrival_airport = 'OVB'  
AND    f.scheduled_departure::date =  
       bookings.now()::date - interval '2 day'  
AND    bp.seat_no = '1A';
```

**Problem.** How many seats were unoccupied on flight PG0404 yesterday?

**Solution.** There are several approaches to solving this problem. One of the options is to use the `NOT EXISTS` expression to find the seats without boarding passes:

```
SELECT count(*)
FROM   flights f
      JOIN seats s
        ON s.aircraft_code = f.aircraft_code
WHERE  f.flight_no = 'PG0404'
AND    f.scheduled_departure::date =
       bookings.now()::date - interval '1 day'
AND    NOT EXISTS (
      SELECT NULL
      FROM   boarding_passes bp
      WHERE  bp.flight_id = f.flight_id
      AND    bp.seat_no = s.seat_no
    );
```

Another approach uses set subtraction to find the unoccupied seats. Different solutions give the same result but may sometimes differ in performance, so you have to take it into account if it matters.

```
SELECT count(*) FROM
(
  SELECT s.seat_no
  FROM   seats s
  WHERE  s.aircraft_code = (
    SELECT aircraft_code FROM flights
    WHERE flight_no = 'PG0404'
    AND    scheduled_departure::date =
           bookings.now()::date - interval '1 day'
  )
  EXCEPT
  SELECT bp.seat_no
  FROM   boarding_passes bp
  WHERE  bp.flight_id = (
    SELECT flight_id FROM flights
    WHERE flight_no = 'PG0404'
    AND    scheduled_departure::date =
           bookings.now()::date - interval '1 day'
  )
) t;
```

76 **Problem.** Which flights had the longest delays? Print the list  
v of ten “leaders.”

**Solution.** The query needs to include only those flights that have already departed.

```
SELECT    f.flight_no,  
          f.scheduled_departure,  
          f.actual_departure,  
          f.actual_departure - f.scheduled_departure  
          AS delay  
FROM      flights f  
WHERE     f.actual_departure IS NOT NULL  
ORDER BY  f.actual_departure - f.scheduled_departure  
          DESC  
LIMIT 10;
```

You can define the same condition using the status column by listing all the applicable statuses. Or you can skip the WHERE condition altogether by specifying the DESC NULLS LAST sorting order, so that undefined values are returned at the end of the selection.

## Aggregate Functions

**Problem.** What is the shortest flight duration for each possible flight from Moscow to St. Petersburg, and how many times was the flight delayed for more than an hour?

**Solution.** To solve this problem, it is convenient to use the available flights\_v view instead of dealing with table joins. You need to take into account only those flights that have already arrived.

```
SELECT    f.flight_no,
          f.scheduled_duration,
          min(f.actual_duration),
          max(f.actual_duration),
          sum(CASE WHEN f.actual_departure >
                      f.scheduled_departure +
                      interval '1 hour'
                THEN 1 ELSE 0
            END) delays
FROM      flights_v f
WHERE     f.departure_city = 'Moscow'
AND       f.arrival_city = 'St. Petersburg'
AND       f.status = 'Arrived'
GROUP BY f.flight_no,
          f.scheduled_duration;
```

**Problem.** Identify the passengers who were the first to check in for all their flights. Consider only those who have taken at least two flights.

**Solution.** Use the fact that boarding pass numbers are issued in the check-in order.

```
SELECT    t.passenger_name,
          t.ticket_no
FROM      tickets t
          JOIN boarding_passes bp
            ON bp.ticket_no = t.ticket_no
GROUP BY t.passenger_name,
          t.ticket_no
HAVING    max(bp.boarding_no) = 1
AND       count(*) > 1;
```

**Problem.** How many passengers can be included into a single booking?

**Solution.** Let's count the number of passengers in each booking and then find the number of bookings for each number of passengers.

```

78  SELECT  tt.cnt,
v    count(*)
    FROM  (
        SELECT  t.book_ref,
                count(*) cnt
        FROM    tickets t
        GROUP BY t.book_ref
    ) tt
GROUP BY tt.cnt
ORDER BY tt.cnt;

```

## Window Functions

**Problem.** For each ticket, display all the included flight segments, together with connection time. Limit the result to the tickets booked a week ago.

**Solution.** Use window functions to avoid accessing the same data twice.

```

SELECT  tf.ticket_no,
        f.departure_airport,
        f.arrival_airport,
        f.scheduled_arrival,
        lead(f.scheduled_departure) OVER w
        AS next_departure,
        lead(f.scheduled_departure) OVER w -
        f.scheduled_arrival AS gap
FROM    bookings b
        JOIN tickets t
        ON t.book_ref = b.book_ref
        JOIN ticket_flights tf
        ON tf.ticket_no = t.ticket_no
        JOIN flights f
        ON tf.flight_id = f.flight_id
WHERE   b.book_date =
        bookings.now()::date - interval '7 day'
WINDOW w AS (PARTITION BY tf.ticket_no
              ORDER BY f.scheduled_departure);

```

As you can see, the time cushion between flights can reach up to several days: round-trip tickets and one-way tickets are treated in the same way, and the time of the stay in the point of destination is treated just like the time between connecting flights. Using the solution for one of the problems in the “Arrays” section, you can take this fact into account when building the query.

**Problem.** What are the most frequent combinations of first and last names? What is the ratio of the passengers with such names to the total number of passengers?

**Solution.** The total number of passengers is calculated using a window function.

```
SELECT passenger_name,  
       round( 100.0 * cnt / sum(cnt) OVER (), 2)  
AS percent  
FROM (   
      SELECT passenger_name,  
             count(*) cnt  
      FROM   tickets  
      GROUP BY passenger_name  
    ) t  
ORDER BY percent DESC;
```

**Problem.** Solve the previous problem for first names and last names separately.

**Solution.** Let's take a look at how to count first names. The query for counting last names will differ only by the p sub-query.

As this complex query shows, you should avoid using a single text field for different values if you are going to use them separately; in scientific terms, it is called the first normal form.

```

80 WITH p AS (
v   SELECT left(passenger_name,
              position(' ' IN passenger_name))
      AS passenger_name
   FROM   tickets
  )
SELECT passenger_name,
       round( 100.0 * cnt / sum(cnt) OVER (), 2)
   AS percent
FROM   (
      SELECT  passenger_name,
              count(*) cnt
        FROM   p
       GROUP BY passenger_name
  ) t
ORDER BY percent DESC;

```

## Arrays

**Problem.** The database does not explicitly indicate whether the ticket is one-way or round-trip. However, you can figure it out by comparing the first point of departure to the last point of destination. Display airports of departure and destination for each ticket, ignoring connections, and specify whether it's a round-trip ticket or not.

**Solution.** An easy solution is to use the `array_agg` function to transform the list of airports in the itinerary into an array.

We select the middle element of the array as the airport of destination, assuming that the outbound and inbound ways have the same number of stops.

In this example, the `tickets` table is scanned only once. The array of airports is displayed for clarity; for large volumes of data, it makes sense to remove it from the query since extra data can hamper performance.

```
WITH t AS (  
  SELECT ticket_no,  
         a,  
         a[1] departure,  
         a[cardinality(a)] last_arrival,  
         a[cardinality(a)/2+1] middle  
  FROM (  
    SELECT t.ticket_no,  
           array_agg( f.departure_airport  
                     ORDER BY f.scheduled_departure) ||  
           (array_agg( f.arrival_airport  
                     ORDER BY f.scheduled_departure DESC)  
            ) [1] AS a  
    FROM tickets t  
         JOIN ticket_flights tf  
           ON tf.ticket_no = t.ticket_no  
         JOIN flights f  
           ON f.flight_id = tf.flight_id  
    GROUP BY t.ticket_no  
  ) t  
)  
SELECT t.ticket_no,  
       t.a,  
       t.departure,  
       CASE  
         WHEN t.departure = t.last_arrival  
          THEN t.middle  
         ELSE t.last_arrival  
       END arrival,  
       (t.departure = t.last_arrival) return_ticket  
FROM t;
```

**Problem.** Find the round-trip tickets in which the outbound route differs from the inbound one.

**Problem.** Find pairs of airports with inbound and outbound flights departing on different days of the week.

**Solution.** The part of the problem that involves building an array of days of the week is virtually solved in the routes

82 view. You only have to find the intersection of arrays using  
v the && operator and make sure it's empty.

```
SELECT r1.departure_airport,  
       r1.arrival_airport,  
       r1.days_of_week dow,  
       r2.days_of_week dow_back  
FROM   routes r1  
       JOIN routes r2  
       ON r1.arrival_airport = r2.departure_airport  
       AND r1.departure_airport = r2.arrival_airport  
WHERE  NOT (r1.days_of_week && r2.days_of_week);
```

## Recursive Queries

**Problem.** How can you get from Ust-Kut (UKX) to Neryungri (CNN) with the minimal number of connections? What will the flight time be?

**Solution.** This problem requires finding the shortest path in a graph. We will use a recursive query to complete this task.

A detailed step-by-step explanation of this query is published at [habr.com/en/company/postgrespro/blog/490228/](http://habr.com/en/company/postgrespro/blog/490228/). Below is a brief summary.

Infinite looping is prevented by checking the hops array, which is built while the query is being executed.

Since the query performs a breadth-first search, the first path found will have the fewest connections. To avoid looping through other paths (that can be numerous and are definitely longer than the already found one), the found attribute is used. It is calculated using the `bool_or` window function.

```
WITH RECURSIVE p(  
    last_arrival, destination, hops,  
    flights, flight_time, found  
) AS (  
    SELECT a_from.airport_code,  
           a_to.airport_code,  
           array[a_from.airport_code],  
           array[]::char(6)[],  
           interval '0',  
           a_from.airport_code = a_to.airport_code  
    FROM   airports a_from,  
           airports a_to  
    WHERE  a_from.airport_code = 'UKX'  
    AND    a_to.airport_code = 'CNN'  
    UNION ALL  
    SELECT r.arrival_airport,  
           p.destination,  
           (p.hops || r.arrival_airport)::char(3)[],  
           (p.flights || r.flight_no)::char(6)[],  
           p.flight_time + r.duration,  
           bool_or(r.arrival_airport = p.destination)  
           OVER ()  
    FROM   p JOIN routes r  
           ON r.departure_airport = p.last_arrival  
    WHERE  NOT r.arrival_airport = ANY(p.hops)  
    AND    NOT p.found  
)  
SELECT hops, flights, flight_time  
FROM   p  
WHERE  p.last_arrival = p.destination;
```

It is useful to compare this query with its simpler variant without the found trick.

To learn more about recursive queries, see the documentation: [postgrespro.com/doc/queries-with](http://postgrespro.com/doc/queries-with).

**Problem.** What is the maximum number of connections that can be required to get from any airport to any other airport?

84 **Solution.** We can take the previous query as the basis for the  
V solution. However, the first iteration must now contain all the possible airport pairs, not just a single pair: each airport must be connected to all the other airports. For all these pairs of airports we first find the shortest path, and then select the longest of them.

Clearly, it is only possible if the routes graph is connected, but our demo database satisfies this condition.

```
WITH RECURSIVE p(  
    departure, last_arrival,  
    destination, hops, found  
) AS (  
    SELECT a_from.airport_code,  
           a_from.airport_code,  
           a_to.airport_code,  
           array[a_from.airport_code],  
           a_from.airport_code = a_to.airport_code  
    FROM   airports a_from,  
           airports a_to  
    UNION ALL  
    SELECT p.departure,  
           r.arrival_airport,  
           p.destination,  
           (p.hops || r.arrival_airport)::char(3)[],  
           bool_or(r.arrival_airport = p.destination)  
           OVER (PARTITION BY p.departure,  
                    p.destination)  
    FROM   p JOIN routes r  
           ON r.departure_airport = p.last_arrival  
    WHERE  NOT r.arrival_airport = ANY(p.hops)  
    AND    NOT p.found  
)  
SELECT max(cardinality(hops)-1)  
FROM   p  
WHERE  p.last_arrival = p.destination;
```

This query also uses the found attribute, but here it should be calculated separately for each pair of airports.

**Problem.** Find the shortest route from Ust-Kut (UKX) to Neryungri (CNN) from the flight time perspective (ignoring connection time).

85  
v

**Solution.** To avoid infinite looping, we use the CYCLE clause introduced in PostgreSQL 14.

```
WITH RECURSIVE p(
  last_arrival, destination,
  flights, flight_time,
  min_time
) AS (
  SELECT a_from.airport_code,
         a_to.airport_code,
         array[]::char(6)[],
         interval '0',
         NULL::interval
  FROM   airports a_from,
         airports a_to
  WHERE  a_from.airport_code = 'UKX'
  AND    a_to.airport_code = 'CNN'
  UNION ALL
  SELECT r.arrival_airport,
         p.destination,
         (p.flights || r.flight_no)::char(6)[],
         p.flight_time + r.duration,
         least(
           p.min_time, min(p.flight_time+r.duration)
         )
  FILTER (
    WHERE r.arrival_airport = p.destination
  ) OVER ()
  FROM   p
  JOIN   routes r
        ON r.departure_airport = p.last_arrival
  WHERE  p.flight_time + r.duration <
         coalesce(p.min_time, interval '1 year')
)
CYCLE last_arrival SET is_cycle USING hops
```

```

86  SELECT hops,
    v      flights,
          flight_time
FROM      (
          SELECT hops,
                 flights,
                 flight_time,
                 min(min_time) OVER () min_time
          FROM    p
          WHERE   p.last_arrival = p.destination
          ) t
WHERE     flight_time = min_time;

```

Note that the found route may be suboptimal with regards to the number of connections.

## Functions and Extensions

**Problem.** Find the distance between Kaliningrad (KGD) and Petropavlovsk-Kamchatsky (PKC).

**Solution.** The airports table contains airport coordinates. To accurately calculate the distance between two points, the Earth's curvature must be considered. This task is best performed by the PostGIS extension, which can approximate the Earth surface as a geoid.

However, a simple spherical model is sufficient for this example. Let's use the earthdistance and then convert the result from miles to kilometers.

```
CREATE EXTENSION IF NOT EXISTS cube;
```

```
CREATE EXTENSION IF NOT EXISTS earthdistance;
```

```
SELECT round(  
    (a_from.coordinates <@> a_to.coordinates) *  
    1.609344  
)  
FROM airports a_from,  
     airports a_to  
WHERE a_from.airport_code = 'KGD'  
AND   a_to.airport_code = 'PKC';
```

87  
v

**Problem.** Draw the graph of flights between all the airports.



# VI PostgreSQL for Applications

## A Separate User

In the previous chapter, we showed how to connect to the database server on behalf of the postgres user. This is the only database user available right after the PostgreSQL installation. Since the postgres user is a superuser, it should not be used to connect to the database from an application. It is better to create a new user and make it the owner of a separate database; then its rights will be limited to this database.

```
postgres=# CREATE USER app PASSWORD 'p@ssw0rd';
CREATE ROLE
postgres=# CREATE DATABASE appdb OWNER app;
CREATE DATABASE
```

To learn about users and privileges, see: [postgrespro.com/doc/user-manag](https://www.postgresql.org/docs/user-manag) and [postgrespro.com/doc/ddl-priv](https://www.postgresql.org/docs/ddl-priv).

To connect to a new database and start working with it on behalf of the newly created user, run:

```
postgres=# \c appdb app localhost 5432
```

```
90 Password for user app: ***
vi You are now connected to database "appdb" as user
   "app" on host "127.0.0.1" at port "5432".

appdb=>
```

This command takes four parameters, in the following order: database name (appdb), username (app), node (localhost or 127.0.0.1), and port number (5432).

Note that the database name is not the only thing that has changed in the prompt: instead of the hash symbol (#), the greater than sign is displayed (>). The hash symbol indicates the superuser rights, similar to the root user in Unix.

The app user has full privileges within the appdb database. For example, this user can create a table:

```
appdb=> CREATE TABLE greeting(s text);
CREATE TABLE
appdb=> INSERT INTO greeting VALUES ('Hello, world!');
INSERT 0 1
```

## Remote Connections

In our example, both the client and the database are located on the same system. You can install PostgreSQL onto a separate server and connect to it from a different system (for example, from an application server). In this case, you must specify your database server address instead of localhost. But it is not enough: for security reasons, PostgreSQL only allows local connections by default.

To connect to the database from the outside, you must edit two files.

First of all, modify the `postgresql.conf` file, which contains **the main configuration settings**. It is usually located in the data directory.

Find the line defining network interfaces for PostgreSQL to listen on:

```
#listen_addresses = 'localhost'
```

We have to replace it with:

```
listen_addresses = '*'
```

Next, edit the `pg_hba.conf` file with **authentication settings**.

When a client tries to connect to the server, PostgreSQL searches this file for the first line that matches the connection by four parameters: connection type, database name, username, and client IP address. This line also specifies how the user must confirm their identity.

For example, on Debian and Ubuntu, this file includes the following setting among others (the top line starting with the hash symbol is a comment):

```
# TYPE DATABASE USER ADDRESS METHOD
local all all peer
```

It means that local connections (`local`) to any database (`all`) on behalf of any user (`all`) must be validated by the peer authentication method (clearly, an IP address is not required for local connections).

92 The peer method means that PostgreSQL requests the current  
vi username from the operating system and assumes that the OS has already performed the required authentication check (prompted for the password). This is why on Linux-like operating systems users usually don't have to enter the password when connecting to a local server.

But Windows does not support local connections, so this line looks as follows:

```
# TYPE DATABASE USER ADDRESS METHOD
host all all 127.0.0.1/32 md5
```

It means that network connections (host) to any database (all) on behalf of any user (all) from the local address (127.0.0.1) must be checked by the md5 method. This method requires the user to enter the password.

To allow the app user to access the appdb database from any address upon providing a valid password, add the following line to the end of the pg\_hba.conf file:

```
host appdb app all md5
```

After changing the configuration files, don't forget to make the server re-read the settings:

```
postgres=# SELECT pg_reload_conf();
```

To learn more details about authentication settings, see [postgrespro.com/doc/client-authentication.html](https://postgrespro.com/doc/client-authentication.html)

To access PostgreSQL from an application, you have to find an appropriate library and install the corresponding database driver. A driver is usually a wrapper for libpq (a standard library that implements the client-server protocol for PostgreSQL), but other implementations are also possible. The library provides application developers with a convenient way to access low-level features of the protocol.

The following sections contain simple code snippets in several popular languages. These examples can help you quickly check the connection with the database system that you have installed and set up.

These code snippets are minimal examples and should not be used in production; there is nothing else, not even error handling functionality.

If you are using a Windows system, to ensure the correct display of extended character sets, you may need to switch to a TrueType font (such as “Lucida Console” or “Consolas”) in the Command Prompt window and change the code page. For example, for the Russian language, run the following commands:

```
C:\> chcp 1251
```

```
Active code page: 1251
```

```
C:\> set PGCLIENTENCODING=WIN1251
```

## PHP

PHP interacts with PostgreSQL via a special extension. On Linux, apart from the PHP itself, you also have to install the package with this extension:

```
$ sudo apt-get install php-cli php-pgsql
```

You can install PHP for Windows from the PHP website: [windows.php.net/download](http://windows.php.net/download). The extension for PostgreSQL is already included in the binary distribution, but you must find and uncomment (by removing the semicolon) the following line in the `php.ini` file:

```
;extension=php_pgsql.dll
```

A sample program (`test.php`):

```
<?php
    $conn = pg_connect('host=localhost port=5432 ' .
                      'dbname=appdb user=app ' .
                      'password=p@ssw0rd') or die;
    $query = pg_query('SELECT * FROM greeting') or die;
    while ($row = pg_fetch_array($query)) {
        echo $row[0].PHP_EOL;
    }
    pg_free_result($query);
    pg_close($conn);
?>
```

Let's execute this command:

```
$ php test.php
Hello, world!
```

You can read about this PostgreSQL extension in PHP documentation: [php.net/manual/en/book.pgsql.php](http://php.net/manual/en/book.pgsql.php).

In the Perl language, database operations are implemented via the DBI interface. On Debian and Ubuntu, Perl itself is pre-installed, so you only need to install the driver:

```
$ sudo apt-get install libdbd-pg-perl
```

There are several Perl builds for Windows, which are listed at [perl.org/get.html](http://perl.org/get.html). Popular builds such as ActiveState Perl and Strawberry Perl come with the PostgreSQL driver pre-installed.

A sample program (test.pl):

```
use DBI;
use open ':std', ':utf8';
my $conn = DBI->connect(
    'dbi:Pg:dbname=appdb;host=localhost;port=5432',
    'app', 'p@ssw0rd') or die;
my $query = $conn->prepare('SELECT * FROM greeting');
$query->execute() or die;
while (my @row = $query->fetchrow_array()) {
    print @row[0]."\n";
}
$query->finish();
$conn->disconnect();
```

Let's execute this command:

```
$ perl test.pl
```

```
Hello, world!
```

The interface is described in the documentation:  
[metacpan.org/pod/DBD::Pg](http://metacpan.org/pod/DBD::Pg).

## Python

Python typically uses the `psycopg` library (pronounced as “psycho-pee-gee”) to interact with PostgreSQL.

On modern versions of Debian and Ubuntu, Python 3 is pre-installed, so you only need to add the corresponding driver:

```
$ sudo apt-get install python3-psycopg2
```

You can download Python for Windows from the [python.org](https://python.org) website. The `psycopg` library is available at [initd.org/psycopg](https://initd.org/psycopg) (choose the version that corresponds to the version of Python installed). You can also find all the required documentation there.

A sample program (`test.py`):

```
import psycopg2
conn = psycopg2.connect(
    host='localhost', port='5432', database='appdb',
    user='app', password='p@ssw0rd')
cur = conn.cursor()
cur.execute('SELECT * FROM greeting')
rows = cur.fetchall()
for row in rows:
    print(row[0])
conn.close()
```

Let's execute this command:

```
$ python3 test.py
```

```
Hello, world!
```

In Java, databases are accessed via the JDBC interface. Java SE 11 is required along with a JDBC driver package.

```
$ sudo apt-get install openjdk-11-jdk
$ sudo apt-get install libpostgresql-jdbc-java
```

You can download JDK for Windows from [oracle.com/technetwork/java/javase/downloads](https://www.oracle.com/technetwork/java/javase/downloads). The JDBC driver is available at [jdbc.postgresql.org](https://jdbc.postgresql.org) (choose the version that corresponds to the JDK installed on your system). You can also find all the required documentation there.

Let's consider a sample program (Test.java):

```
import java.sql.*;
public class Test {
    public static void main(String[] args)
        throws SQLException {
        Connection conn = DriverManager.getConnection(
            "jdbc:postgresql://localhost:5432/appdb",
            "app", "p@ssw0rd");
        Statement st = conn.createStatement();
        ResultSet rs = st.executeQuery(
            "SELECT * FROM greeting");
        while (rs.next()) {
            System.out.println(rs.getString(1));
        }
        rs.close();
        st.close();
        conn.close();
    }
}
```

Compile and execute the program specifying the path to the JDBC class driver (on Windows, paths are separated by semi-colons, not colons):

```
98 $ javac Test.java
vi  $ java -cp ./usr/share/java/postgresql-jdbc4.jar \
    Test
    Hello, world!
```

## Backup

Although our database contains only one table, ensuring data durability is still important. While your application has little data, the easiest way to create a backup is to use the `pg_dump` utility:

```
$ pg_dump appdb > appdb.dump
```

If you open the resulting `appdb.dump` file in a text editor, you will see standard SQL commands that create all the `appdb` objects and fill them with data. You can pass this file to `psql` to restore the contents of the database. For example, you can create a new database and import all the data into it:

```
$ createdb appdb2
$ psql -d appdb2 -f appdb.dump
```

This is the format in which the demo database described in the previous chapter is distributed.

The `pg_dump` utility offers many features worth checking out: [postgrespro.com/doc/app-pgdump](https://www.postgresql.org/docs/app-pgdump). Some of them are available only if the data is dumped in a custom format. In this case, you have to use the `pg_restore` utility instead of `psql` to restore the data.

In any case, `pg_dump` can back up the contents of a single database only. To make a backup of the whole cluster, including all the databases, users, and tablespaces, you should use a different command: `pg_dumpall`.

Large-scale projects require an elaborate and comprehensive backup strategy. A better option here is a physical binary copy of the cluster, which can be taken by the `pg_basebackup` utility:

```
$ pg_basebackup -D backup
```

This command will create a backup of the whole database cluster in the backup directory. To restore the cluster from this backup, move it to the data directory and restart the server.

To learn more about the available backup and recovery tools, see the documentation: [postgrespro.com/doc/backup](https://www.postgresql.org/docs/current/pgbasebackup.html).

Built-in PostgreSQL features enable you to implement almost everything you need, but you have to complete multi-step workflows that lack automation. That's why many companies create their own backup tools. Postgres Professional also has such a tool called `pg_probackup`. The free version of the tool enables you to perform incremental backups at the page level, ensure data integrity, use parallel execution and compression when working with big volumes of information, and implement various backup strategies.

Its full documentation is available at [postgrespro.com/doc/app-pgprobackup](https://www.postgresql.org/docs/current/pgprobackup.html).

## What's next?

Now you are ready to develop your application. Regarding the database, the application will always consist of two parts: server and client. The server part comprises everything that relates to the database system: tables, indexes, views, triggers, stored functions and procedures. The client part holds everything that works outside of the database and connects to it; from the database point of view, it doesn't matter whether it's a thick client or an application server.

An important question that has no clear-cut answer: where should we place business logic?

One popular approach is to move the logic out of the database and implement it all on the client. It often happens when developers are unfamiliar with all the capabilities provided by a relational database system and prefer to rely on what they know well, that is, on the application code.

In this case, the database becomes somewhat secondary to the application and only ensures data persistence, its reliable storage. Database systems can be often isolated by an additional abstraction level, such as an ORM tool that automatically generates database queries from the constructs of the programming language familiar to developers. Such solutions are sometimes justified by the intent to develop an application that is portable to any database system.

This approach has the right to exist: if such a system works and addresses all business objectives, why not?

However, this solution also has some obvious drawbacks:

- **Data consistency is ensured by the application.**  
Instead of relying on the database system to ensure data

consistency (and this is exactly what relational database systems are especially good at), all the required checks are performed by the application. Rest assured that sooner or later your database will contain inconsistent data. You have to either fix these errors, or teach the application how to handle them. If multiple applications use the same database, enforcing data consistency without database-level constraints becomes impractical.

- **Performance leaves much to be desired.**

ORM systems allow you to create an abstraction level over the database, but the quality of SQL queries they generate is rather questionable. As a rule, multiple small queries are executed, and each of them is quite fast on its own. Such a model can cope only with low load on small data volumes and is virtually impossible to optimize on the database side.

- **Application code gets more complicated.**

Using application-oriented programming languages, it's impossible to write a really complex query that could be properly translated to SQL in an automated way. Thus, complex data processing (if it is needed, of course) has to be implemented at the application level, with all the required data retrieved from the database in advance, but it requires an extra data transfer over the network. Furthermore, such algorithms as scans, joins, sorting, and aggregation provided by database systems are guaranteed to perform better than the application code, as they have been improved and optimized for years.

To fully utilize database features, including integrity constraints and stored procedures, a thorough understanding of its capabilities is necessary. You have to master the SQL language to write queries and learn one of the server program-

102      ming languages (typically, PL/pgSQL) to create functions and  
vi      triggers. In return, you will get a reliable tool, one of the  
    most important building blocks for any information system  
    architecture.

In any case, you have to decide for yourself where to implement business logic: on the server side or on the client side. We'll just note that there's no need to go to extremes, as the truth often lies somewhere in the middle.

# VII Configuring PostgreSQL

## Basic Settings

The default settings allow us to start PostgreSQL on virtually any hardware, including low-end systems. But for best performance, the database configuration has to take into account both physical characteristics of the server and a typical application workload.

Here we'll cover only some of the basic settings that must be considered for a production-level database system. Fine-tuning for a specific application requires additional expertise, which can be acquired through PostgreSQL database administration courses (see p. 149).

## Changing Configuration Parameters

To change a configuration parameter, you have to open the `postgresql.conf` file and either find the required parameter and modify its value, or add a new line at the end of the file: it will have priority over the setting specified above in the same file.

After changing the settings, you have to reload the server configuration:

```
104 postgres=# SELECT pg_reload_conf();  
vii
```

Now check the current setting using the `SHOW` command. If the parameter value has not changed, take a look into the server log: you might have made a mistake when editing the file.

Instead of manually editing the configuration file, you can use an SQL command to update the parameter. However, this approach still requires reloading the server configuration:

```
postgres=# ALTER SYSTEM SET work_mem='128MB';
```

Such settings get into the `postgresql.auto.conf` file and take priority over the values specified in the main file. The advantage of this method is that the new parameter values get validated at once.

## The Most Important Parameters

It is highly important to pay attention to parameters that define how PostgreSQL uses RAM.

The `shared_buffers` parameter sets the size of PostgreSQL's shared memory buffers, which store frequently accessed data to minimize disk reads. A reasonable starting value is 25 % of all the RAM used by the server. Changing this parameter requires a server restart.

The `effective_cache_size` value has no effect on memory allocation; it merely prompts the size of cache PostgreSQL can count on, including the operating system cache. The larger the value, the higher priority is given to indexes. You can start with 50–75 % of RAM.

The *work\_mem* parameter defines the amount of memory allocated for sorting, building hash tables when performing joins, and other operations. Frequent use of temporary files suggests insufficient memory allocation, which can degrade performance. In most cases, the default value of 4 MB should be increased by at least several times, but be cautious not to exceed the overall RAM size of the server.

The *maintenance\_work\_mem* parameter defines the amount of memory allocated for service processes. Higher values can speed up indexing and vacuuming. This parameter is usually set to a value that is several times higher than *work\_mem*.

For example, for 32 GB of RAM, you can start with the following settings:

```
shared_buffers = '8GB'  
effective_cache_size = '24GB'  
work_mem = '128MB'  
maintenance_work_mem = '512MB'
```

The ratio of *random\_page\_cost* to *seq\_page\_cost* must match the ratio of random disk access speed to sequential access speed. By default, random access is assumed to be four times slower than sequential one (which works well for regular HDDs). For disk arrays and SSDs, lowering the *random\_page\_cost* value can improve performance. However, the *seq\_page\_cost* value should remain at 1.

For example, the following setting is appropriate for SSD drives:

```
random_page_cost = 1.2
```

It's very important to configure autovacuum. This process performs "garbage collection" and several other critical system

106 tasks. This setting highly depends on a particular application  
vii and its workload.

In most cases, you can start with the following configuration:

- Reduce the *autovacuum\_vacuum\_scale\_factor* value to 0.01 to perform autovacuum more often and in smaller batches.
- Increase the *autovacuum\_vacuum\_cost\_limit* value (or reduce *autovacuum\_vacuum\_cost\_delay*) by 10 times to speed up autovacuum (for version 11 or lower).

It's equally important to properly configure the processes related to buffer cache and WAL maintenance, but the exact settings also depend on a particular application. For a start, you can set the *checkpoint\_completion\_target* parameter to 0.9 (to spread out the load), increase *checkpoint\_timeout* from 5 to 30 minutes (to reduce the overhead caused by checkpoints), and proportionally increase the *max\_wal\_size* value (for the same purpose).

To learn tips and tricks for configuring various parameters, you can take the DBA2 course (p. 154).

## Connection Settings

As discussed in the “PostgreSQL for Applications” chapter (p. 89), the *listen\_addresses* parameter should typically be set to `'*'`, and the `pg_hba.conf` configuration file must be updated to allow connections.

## Bad Advice

You can sometimes find advice about improving performance that should never be followed:

- Turning off autovacuum.

Such “resource saving” may provide minor short-term performance benefits but will result in data fragmentation and bloated tables and indexes. Sooner or later your database system is sure to stop functioning normally. Auto-vacuum should never be turned off, it should be properly configured.

- Turning off disk synchronization (*fsync = off*).

Disabling *fsync* will indeed bring a tangible performance improvement, but any server crash (caused by either software or hardware failure) will lead to a complete loss of all databases. In this case, you can only restore the system from a backup (if you happen to have one).

## PostgreSQL and 1C Solutions

PostgreSQL is officially supported by 1C, a widely used Russian ERP system. It provides a cost-effective alternative to expensive commercial database licenses.

As any other applications, 1C products will work much faster if PostgreSQL is configured appropriately. Additionally, specific server parameters are required for optimal performance with 1C.

Here we'll provide some installation and setup instructions that can help you get started.

## Choosing PostgreSQL Version

1C requires a custom patched version of PostgreSQL. You can download it from [releases.1c.ru](https://releases.1c.ru), or use Postgres Pro Standard or Postgres Pro Enterprise, which also include all the required patches.

PostgreSQL can work on Windows as well, but if you have a choice, it's better to opt for a Linux distribution.

Before you start the installation, you have to decide whether a dedicated database server is required. A dedicated server offers higher performance because of better load balancing between the application server and the database server.

## Configuration Parameters

Physical specifications of the server must match the expected load. You can use the following data as a baseline: a dedicated 8-core server with 8 GB of RAM and a disk subsystem with RAID1 SSD should be enough for a database of 100 GB, the total number of 50 users, and up to 2000 documents per day. If the server is not dedicated, PostgreSQL must get the corresponding amount of resources from the common server.

Based on the general recommendations and 1C application requirements, the following initial settings are recommended for such a server:

```
# Mandatory settings for 1C
standard_conforming_strings = off
escape_string_warning = off
shared_preload_libraries = 'online_analyze, plantuner'
plantuner.fix_empty_table = on
```

```
online_analyze.enable = on
online_analyze.table_type = 'temporary'
online_analyze.local_tracking = on
online_analyze.verbose = off

# The following settings depend on the available RAM
shared_buffers = '2GB'           # 25% of RAM
effective_cache_size = '6GB'     # 75% of RAM
work_mem = '64MB'                # 64-128MB
maintenance_work_mem = '256MB'  # 4*work_mem

# Active use of temporary tables
temp_buffers = '32MB'           # 32-128MB

# The default value of 64 is not enough
max_locks_per_transaction = 256
```

## Connection Settings

Ensure that the *listen\_addresses* parameter in the `postgresql.conf` file is set to `*`.

Add the following line at the start of the `pg_hba.conf` configuration file, specifying the actual address and subnet mask instead of the “IP-address-of-the-1C-server” placeholder:

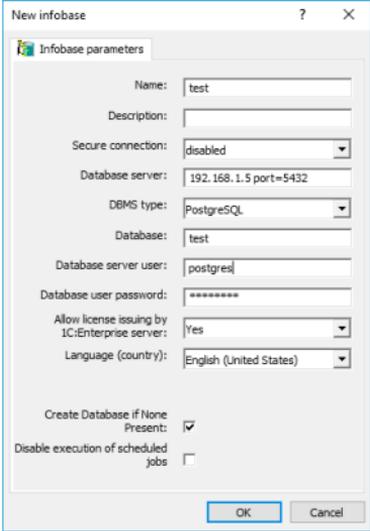
```
host all all IP-address-of-1C-server md5
```

After restarting PostgreSQL, all the changes in `pg_hba.conf` and `postgresql.conf` files take effect, and the server is ready to accept 1C connections.

1C establishes a connection as a superuser, usually `postgres`. Set a password for this role:

```
postgres=# ALTER ROLE postgres PASSWORD 'p@ssw0rd';
ALTER ROLE
```

110 In configuration settings  
vii of the 1C information  
database, specify the  
IP-address and port of  
the PostgreSQL server  
as your database server  
and choose “PostgreSQL”  
as the required DBMS  
type. Specify the name  
of the database that will  
be used for 1C and select  
the “Create database if  
none present” check box  
(do not create this  
database using  
PostgreSQL means).  
Provide the name and  
password of a superuser  
role that will be used  
to establish connections.



The screenshot shows a dialog box titled "New infobase" with a standard Windows window border. The main area is titled "Infobase parameters" and contains the following fields and options:

- Name: test
- Description: (empty)
- Secure connection: disabled (dropdown menu)
- Database server: 192.168.1.5 port=5432
- DBMS type: PostgreSQL (dropdown menu)
- Database: test
- Database server user: postgres
- Database user password: (masked with asterisks)
- Allow license issuing by 1C:Enterprise server: Yes (dropdown menu)
- Language (country): English (United States) (dropdown menu)
- Create Database if None Present:
- Disable execution of scheduled jobs:

At the bottom right, there are "OK" and "Cancel" buttons.

These recommendations should help you to quickly get started, even though they cannot guarantee optimal performance.

We thank Anton Doroshkevich from the Infrosoft company for his assistance in preparing these recommendations.

# VIII pgAdmin

pgAdmin is a widely used graphical tool for managing PostgreSQL databases. It simplifies common database administration tasks, provides an interface for exploring database objects, and allows users to execute SQL queries.

For many years, pgAdmin 3 was the de facto standard for PostgreSQL administration. However, in 2016, EnterpriseDB discontinued support for pgAdmin 3 and introduced pgAdmin 4, which was completely rewritten in Python and modern web technologies, replacing the original C++ implementation. Initially, pgAdmin 4 received mixed reviews due to its redesigned interface. However, ongoing development has led to significant improvements, making it a more robust and feature-rich tool.

## Installation

To launch pgAdmin 4 on Windows, use the installer available at [pgadmin.org/download](https://pgadmin.org/download). The installation procedure is straightforward, and the default options are typically sufficient.

For Debian and Ubuntu, add the PostgreSQL repository (see p. 30), then install pgAdmin 4 using the following command:

112 \$ **sudo apt-get** install pgadmin4  
viii

“pgAdmin4” appears in the list of available programs.

The user interface of this program is fully localized into a dozen languages. To switch to another language, click **Configure pgAdmin**, select **Miscellaneous > User language** in the settings window, and then reload the page in your web browser.

## Connecting to a Server

To begin, set up a connection to the server. Click the **Add New Server** button and enter an arbitrary connection name in the **General** tab of the opened window.

In the **Connection** tab, enter hostname or address, port number, username, and password.

If you don't want to enter the password every time, select the **Save password** checkbox. pgAdmin encrypts stored passwords using a master password, which you will be prompted to set upon first use.

Note that this user must already have a password. For example, for the postgres user, you can do it with the following command:

```
postgres=# ALTER ROLE postgres PASSWORD 'p@ssw0rd';
```

Click **Save** to verify the server connection and add it to pgAdmin.

The screenshot shows a dialog box titled "Register - Server" with a close button (X) in the top right corner. The dialog has five tabs: "General", "Connection" (which is selected), "Parameters", "SSH Tunnel", and "Advanced". The "Connection" tab contains the following fields and controls:

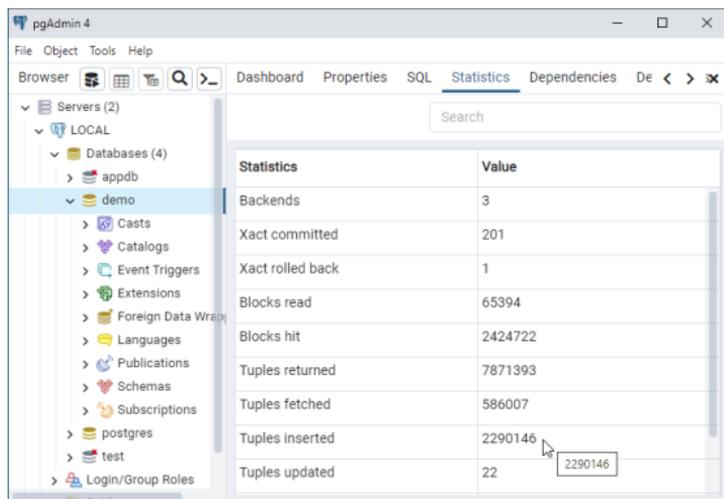
- Host name/address: localhost
- Port: 5432
- Maintenance database: postgres
- Username: postgres
- Kerberos authentication?:
- Password: .....
- Save password?:
- Role: (empty text box)
- Service: (empty text box)

At the bottom of the dialog, there are three buttons: "Close" (with an X icon), "Reset" (with a circular arrow icon), and "Save" (with a floppy disk icon). There are also two small icons on the left: an information icon (i) and a help icon (?).

## Browser

The left pane displays the Browser tree. Expanding the tree reveals the server, labeled LOCAL in this example. At the next level, you can see all its databases:

- appdb has been created to check connection to PostgreSQL using different programming languages.
- demo is our demo database.
- postgres is always created when PostgreSQL gets installed.
- test was used in the “Trying SQL” chapter (p. 35).



Expanding the **Schemas** section under `appdb` reveals the `greetings` table, along with its columns, constraints, indexes, and triggers.

The context (right-click) menu provides various actions for each object type, including exporting, importing, assigning privileges, and deleting.

The right pane includes several tabs that display reference information:

- **Dashboard** provides system activity charts.
- **Properties** displays the properties of the object selected in the Browser (data types for columns, etc.)
- **SQL** shows the SQL command used to create the selected object.

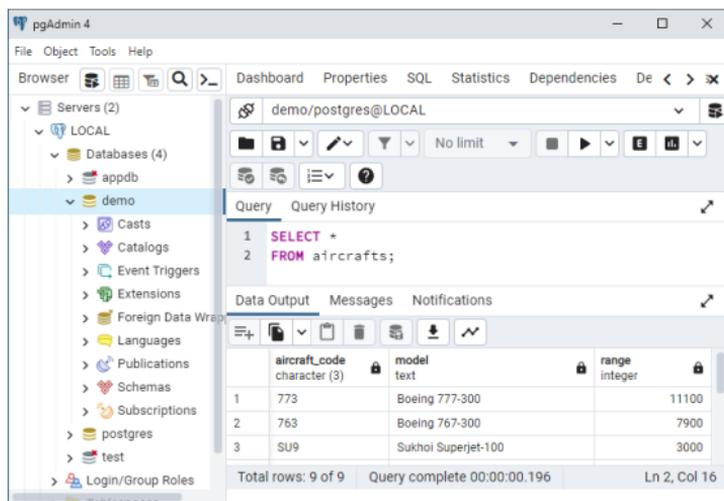
- **Statistics** lists information used by the query optimizer to build query plans; it can be used by a database administrator for case analysis.
- **Dependencies, Dependents** illustrates dependencies between the selected object and other objects in the database.

115  
viii

## Running Queries

To execute a query, open a new SQL editor tab by selecting **Tools > Query Tool** from the menu.

Type your query in the editor and press F5 to execute it. The **Data Output** tab in the bottom panel will display the result of the query.



The screenshot shows the pgAdmin 4 interface. The left sidebar shows a tree view with 'Servers (2)' expanded to 'LOCAL', then 'Databases (4)', and finally 'demo' selected. The main window shows a SQL editor with the following query:

```
1 SELECT *
2 FROM aircrafts;
```

The 'Data Output' tab is active, displaying the results of the query in a table format:

	aircraft_code	model	range
	character (3)	text	integer
1	773	Boeing 777-300	11100
2	763	Boeing 767-300	7900
3	SU9	Sukhoi Superjet-100	3000

At the bottom of the table, it shows 'Total rows: 9 of 9' and 'Query complete 00:00:00.196 Ln 2, Col 16'.

116 To execute a new query without deleting the previous one,  
viii highlight the desired portion of SQL before pressing F5. Thus,  
the whole history of your actions will be always in front of  
you. It is usually more convenient than searching for the  
required query in the log on the **Query History** tab.

## Other Features

pgAdmin provides a graphical user interface for standard PostgreSQL utilities, system catalog information, administration functions, and SQL commands. pgAdmin includes a built-in PL/pgSQL debugger for troubleshooting stored procedures. You can learn about pgAdmin features either on the product website [pgadmin.org](http://pgadmin.org) or in the built-in pgAdmin help system.

# IX Additional Features

## Full-Text Search

Full-text search refers to searching for documents written in natural language and sorting the results by relevance to the search query. In the simplest and most typical case, the query consists of one or more words, and the relevance is defined by the frequency of these words in the document. This is more or less what happens when we type a phrase in Google or Yandex search engines. However, despite SQL's powerful capabilities, it is not always sufficient for effective data handling. This has become increasingly evident with the rise of Big Data, which is often poorly structured and difficult to parse.

There is a large number of search engines, free and paid, that enable you to index the whole collection of your documents and set up search of quite decent quality. In this case, the index—an essential search tool for speeding up queries—is separate from the database. It means that many valuable database features become unavailable: database synchronization, transaction isolation, accessing and using metadata to limit the search range, setting up access policies, and many more.

Shortcomings of document-oriented database management systems usually have a similar nature: they have rich full-text search functionality, but data security and synchronization

118 features are of low priority. Besides, such databases (for ex-  
ix ample, MongoDB) are usually NoSQL ones, so by design, they  
lack all the power of SQL accumulated over years.

However, traditional SQL database systems do have built-in full-text search capabilities. The LIKE operator is part of the standard SQL syntax, but its flexibility is often insufficient. Therefore, developers had to implement their own extensions of the SQL standard. In PostgreSQL, these are comparison operators ILIKE, ~, ~\*, but they don't solve all the problems either, as they don't take into account grammatical forms, are not suitable for ranking, and work rather slowly.

When talking about the tools of the full-text search itself, it's important to understand that they are far from being standardized: each database system uses its own approach and syntax. Russian users of PostgreSQL have some advantage here: its full-text search extensions were created by Russian developers, so there is a possibility to contact the experts directly or even attend their lectures to go into low-level details, if required. Here we'll only provide some simple examples.

To learn about the full-text search capabilities, we are going to create one more table in the demo database. Let it be a lecturer's draft notes split into chapters by lecture topics:

```
test=# CREATE TABLE course_chapters(  
    c_no text REFERENCES courses(c_no),  
    ch_no text,  
    ch_title text,  
    txt text,  
    CONSTRAINT pkt_ch PRIMARY KEY(ch_no, c_no)  
);  
  
CREATE TABLE
```

Now let's enter the text of the first lectures for our courses CS301 and CS305:

119  
ix

```
test=# INSERT INTO course_chapters(
  c_no, ch_no, ch_title, txt)
VALUES
('CS301', 'I', 'Databases',
 'We start our acquaintance with ' ||
 'the fascinating world of databases'),
('CS301', 'II', 'First Steps',
 'Getting more fascinated with ' ||
 'the world of databases'),
('CS305', 'I', 'Local Networks',
 'Here we start our adventurous journey ' ||
 'through the intriguing world of networks');
INSERT 0 3
```

Check the result:

```
test=# SELECT ch_no AS no, ch_title, txt
FROM course_chapters \gx
-[ RECORD 1 ]-----
no      | I
ch_title | Databases
txt      | We start our acquaintance with the
         | fascinating world of databases
-[ RECORD 2 ]-----
no      | II
ch_title | First Steps
txt      | Getting more fascinated with the world of
         | databases
-[ RECORD 3 ]-----
no      | I
ch_title | Local Networks
txt      | Here we start our adventurous journey
         | through the intriguing world of networks
```

Now let's find some information in our database with the help of traditional SQL means (using the LIKE operator):

```
120 test=# SELECT ch_no AS no, ch_title, txt
ix FROM course_chapters
WHERE txt LIKE '%fascination%' \gx
```

It's easy to guess the result: 0 rows. The LIKE operator sees no connection between “fascination” and the words “fascinating” and “fascinated” present in the text.

The query

```
test=# SELECT ch_no AS no, ch_title, txt
FROM course_chapters
WHERE txt LIKE '%fascinated%' \gx
```

will return the row from chapter II (but not from chapter I, where the adjective “fascinating” is used):

```
-[ RECORD 1 ]-----
no          | II
ch_title    | First Steps
txt         | Getting more fascinated with the world of
           | databases
```

PostgreSQL provides the ILIKE operator, which allows us not to worry about letter cases; otherwise, you would also have to take uppercase and lowercase letters into account. True, there are also regular expressions (search patterns), and setting them up is a truly engaging task, little short of art, but sometimes you just want a tool that can simply do the job. So let's add an additional column to the `course_chapters` table; it will have a special type called `tsvector`:

```
test=# ALTER TABLE course_chapters
      ADD txtvector tsvector;
test=# UPDATE course_chapters
      SET txtvector = to_tsvector('english',txt);
test=# SELECT txtvector
FROM course_chapters \gx
```

```

-[ RECORD 1 ]-----
txtvector | 'acquaint':4 'databas':10 'fascin':7
          | 'start':2 'world':9
-[ RECORD 2 ]-----
txtvector | 'databas':8 'fascin':3 'get':1 'world':6
-[ RECORD 3 ]-----
txtvector | 'adventur':5 'intrigu':9 'journey':6
          | 'network':12 'start':3 'world':10

```

As we can see, the rows have changed:

- Words are reduced to their root forms (lexemes).
- Numbers have appeared. They indicate the word position in the text.
- There are no prepositions included (and neither there would be any conjunctions or other parts of the sentence that are unimportant for search; they are the so-called stop words).

The search will be even more efficient if it includes chapter titles, which are also given more weight in respect to the rest of the text (it can be done using the `setweight` function).

Let's modify the table:

```

test=# UPDATE course_chapters
      SET txtvector =
          setweight(to_tsvector('english',ch_title),'B')
          || ' ' ||
          setweight(to_tsvector('english',txt),'D');

```

```
UPDATE 3
```

```

test=# SELECT txtvector FROM course_chapters \gx
-[ RECORD 1 ]-----
txtvector | 'acquaint':5 'databas':1B,11 'fascin':8
          | 'start':3 'world':10
-[ RECORD 2 ]-----
txtvector | 'databas':10 'fascin':5 'first':1B 'get':3
          | 'step':2B 'world':8

```

```
122  -[ RECORD 3 ]-----  
ix    txtvector | 'intrigu':10 'journey':7 'local':1B  
      'network':2B,13 'start':5 'world':11
```

Lexemes are assigned relative weight markers: B and D (possible options are A, B, C, D). We'll assign real weights when writing queries, which will give us more flexibility.

Fully armed, let's return to search. The `to_tsquery` function resembles the `to_tsvector` function we have seen above: it converts a string to the `tsquery` data type used in queries.

```
test=# SELECT ch_title  
FROM course_chapters  
WHERE txtvector @@  
      to_tsquery('english','fascination & database');  
 ch_title  
-----  
 Databases  
 First Steps  
( rows)
```

You can check that the query `'fascinated & database'` and its other grammatical variants will return the same result. Here we have used the comparison operator `@@`, which plays the same role in full-text search as the `LIKE` operator does in regular search. The syntax of the `@@` operator does not allow natural language expressions with spaces, so words in the query are connected by the “and” logical operator.

The `english` argument indicates the configuration used by PostgreSQL. The configuration defines pluggable dictionaries and the parser, which splits the phrase into separate lexemes.

Despite their name, dictionaries enable all kinds of lexeme transformations. For example, a simple stemmer dictionary

like snowball, which is used by default, reduces the word to its unchangeable part; it allows search to ignore word endings in queries. You can also plug in other dictionaries, for example:

- regular dictionaries like `ispell`, `myspell`, or `hunspell`, which can better handle word morphology
- dictionaries of synonyms
- thesaurus
- `unaccent`, which can remove diacritics from letters

Thanks to assigned weights, the displayed search results are ranked:

```
test=# SELECT ch_title,  
             ts_rank_cd('{0.1, 0.0, 1.0, 0.0}', txtvector, q)  
FROM course_chapters,  
     to_tsquery('english', 'Databases') q  
WHERE txtvector @@ q  
ORDER BY ts_rank_cd DESC;
```

ch_title	ts_rank_cd
Databases	1.1
First Steps	0.1

( rows)

The `{0.1, 0.0, 1.0, 0.0}` array sets the weights. It is an optional argument of the `ts_rank_cd` function. By default, array `{0.1, 0.2, 0.4, 1.0}` corresponds to D, C, B, A. The word's weight affects ranking of the returned row.

In the final experiment, let's modify the display format. Suppose we would like to highlight the found words in the html page using the bold type. The `ts_headline` function specifies the symbols used to highlight words and sets the minimum and maximum number of words displayed in a single line:

```

124 test=# SELECT ts_headline(
ix    'english',
      txt,
      to_tsquery('english', 'world'),
      'StartSel=<b>, StopSel=</b>',
      MaxWords=50, MinWords=5'
)
FROM course_chapters
WHERE to_tsvector('english', txt) @@
      to_tsquery('english', 'world');

-[ RECORD 1 ]-----
ts_headline | with the fascinating <b>world</b> of
             | databases
-[ RECORD 2 ]-----
ts_headline | with the <b>world</b> of databases
-[ RECORD 3 ]-----
ts_headline | through the intriguing <b>world</b> of
             | networks

```

To speed up full-text search, special indexes are used: GiST, GIN, and RUM, which are different from regular database indexes. However, like many other useful full-text search features, they fall outside the scope of this guide.

To learn more about full-text search, see the documentation: [www.postgrespro.com/doc/textsearch](http://www.postgrespro.com/doc/textsearch).

## Using JSON and JSONB

SQL-based relational databases were originally designed with a strong emphasis on data consistency and security, at a time when data volumes were far smaller than they are today. When NoSQL databases appeared, it caused concern in the community: lack of strict consistency support and a much

simpler data structure (at first, it was simply a storage of key–value pairs) provided a remarkable search speedup. Actively using parallel computations, they could process unprecedented volumes of information and were easy to scale.

Once the initial shock had passed, it became clear that such a simple structure was insufficient for most practical tasks. Composite keys were introduced, and then groups of keys appeared. To remain competitive, relational database systems began incorporating features commonly associated with NoSQL.

Since changing the database schema in a relational database incurs high costs, a new JSON data type came in handy. Having a hierarchical structure similar to XML, it was first targeted at JavaScript development (hence JS in the title), including AJAX application development. JSON's flexibility allows developers to store and manage data with dynamic, unpredictable structures without requiring database schema modifications.

Suppose we need to enter personal data into the students demo database. We conducted a survey and collected the information from professors. Some questions in the questionnaire are optional, while other questions include the “add more information at your discretion” and “other” fields. If we followed the traditional approach, the information that does not fit the current structure would require adding multiple new tables and columns with lots of empty fields. As the dataset grows, the whole database may have to be refactored.

We can solve this problem using `json` or `jsonb` types. The `jsonb` type, which appeared after `json`, stores data in a compact binary form and, unlike `json`, supports indexes, which can sometimes speed up search by an order of magnitude.

126 Let's create a table with JSON objects:

ix

```
test=# CREATE TABLE student_details(  
    de_id int,  
    s_id int REFERENCES students(s_id),  
    details json,  
    CONSTRAINT pk_d PRIMARY KEY(s_id, de_id)  
);  
  
test=# INSERT INTO student_details  
    (de_id, s_id, details)  
VALUES  
(1, 1451,  
'{ "merits": "none",  
    "flaws":  
    "immoderate ice cream consumption",  
    "status" : "expelled"  
}' ),  
(2, 1432,  
'{ "hobbies":  
    { "guitarist":  
        { "band": "Postgressors",  
          "guitars":["Strat","Telec"]  
        }  
    }  
}' ),  
(3, 1556,  
'{ "hobbies": "cosplay",  
    "merits":  
    { "mother-of-five":  
        { "Basil": "m", "Simon": "m", "Lucie": "f",  
          "Mark": "m", "Alex": "unknown"  
        }  
    }  
}' );
```

Let's verify that the data has been correctly inserted. For convenience, we will join the `student_details` and `students` tables using the `WHERE` clause, as the new table does not contain students' names:

```
test=# SELECT s.name, sd.details
FROM student_details sd, students s
WHERE s.s_id = sd.s_id \gx

-[ RECORD 1 ]-----
name      | Anna
details   | { "merits": "none",           +
      |   "flaws":                    +
      |     "immoderate ice cream consumption", +
      |     "status" : "expelled"     +
      | }
-[ RECORD 2 ]-----
name      | Victor
details   | { "hobbies":                  +
      |   { "guitarist":             +
      |     { "band": "Postgressors", +
      |       "guitars":["Strat","Telec"] +
      |     }
      |   }
      | }
-[ RECORD 3 ]-----
name      | Nina
details   | { "hobbies": "cosplay",      +
      |   "merits":                  +
      |     { "mother-of-five":     +
      |       { "Basil": "m",      +
      |         "Simon": "m",      +
      |         "Lucie": "f",      +
      |         "Mark": "m",      +
      |         "Alex": "unknown"  +
      |       }
      |     }
      | }
```

Suppose we are interested in entries that hold information about the students' merits. Let's access the values of the "merits" key using a special operator ->>:

```
test=# SELECT s.name, sd.details
FROM student_details sd, students s
WHERE s.s_id = sd.s_id
AND sd.details ->> 'merits' IS NOT NULL \gx
```

```

128  ix  -[ RECORD 1 ]-----
      name      | Anna
      details   | { "merits": "none",
      |           | "flaws":
      |           | "immoderate ice cream consumption",
      |           | "status" : "expelled"
      |           | }
      -[ RECORD 2 ]-----
      name      | Nina
      details   | { "hobbies": "cosplay",
      |           | "merits":
      |           | { "mother-of-five":
      |           |   { "Basil": "m",
      |           |     "Simon": "m",
      |           |     "Lucie": "f",
      |           |     "Mark": "m",
      |           |     "Alex": "unknown"
      |           |   }
      |           | }
      |           | }

```

We have seen that the two entries are related to Anna and Nina's merits, but the result is misleading, as Anna's merits are actually "none." Let's fix the query:

```

test=# SELECT s.name, sd.details
FROM student_details sd, students s
WHERE s.s_id = sd.s_id
AND sd.details ->> 'merits' IS NOT NULL
AND sd.details ->> 'merits' != 'none';

```

The new query only returns Nina, whose merits are real.

However, this method is not always effective. Let's attempt to retrieve the guitars that Victor plays:

```

test=# SELECT sd.de_id, s.name, sd.details
FROM student_details sd, students s
WHERE s.s_id = sd.s_id
AND sd.details ->> 'guitars' IS NOT NULL \gx

```

This query won't return anything. It's because the corresponding key-value pair is located inside the JSON hierarchy, nested into the pairs of a higher level:

129  
ix

```
name      | Victor
details   | { "hobbies":                                +
          |   { "guitarist":                            +
          |     { "band": "Postgressors",              +
          |       "guitars":["Strat","Telec"]         +
          |     }                                       +
          |   }                                         +
          | }                                           +
          | }
```

To retrieve the list of guitars, let's use the #> operator and go down the hierarchy, starting from "hobbies":

```
test=# SELECT sd.de_id, s.name,
           sd.details #> '{hobbies,guitarist,guitars}'
FROM   student_details sd, students s
WHERE  s.s_id = sd.s_id
AND    sd.details #> '{hobbies,guitarist,guitars}'
       IS NOT NULL;
```

We can see that Victor is a fan of Fender:

```
de_id | name | ?column?
-----+-----+-----
      | Victor | ["Strat","Telec"]
```

The json type has a younger sibling: jsonb. The letter "b" implies the binary (rather than text) format of data storage and structure, which can often result in faster search. Nowadays, jsonb is used much more frequently than json.

```
test=# ALTER TABLE student_details
ADD details_b jsonb;
```

```

130 test=# UPDATE student_details
ix   SET details_b = to_jsonb(details);
test=# SELECT de_id, details_b
FROM student_details \gx
-[ RECORD 1 ]-----
de_id      | 1
details_b  | {"flaws": "immoderate ice cream
consumption", "merits": "none",
"status": "expelled"}
-[ RECORD 2 ]-----
de_id      | 2
details_b  | {"hobbies": {"guitarist": {"guitars":
["Strat", "Telec"], "band":
"Postgressors"}}}
-[ RECORD 3 ]-----
de_id      | 3
details_b  | {"hobbies": "cosplay", "merits":
{"mother-of-five": {"Basil": "m",
"Lucie": "f", "Alex": "unknown",
"Mark": "m", "Simon": "m"}}}

```

Besides the notation differences, the order of values has changed: Alex, who lacks additional information, now appears before Mark. It's not a disadvantage of jsonb as compared to json, it's simply its data storage specifics.

The jsonb type is supported by a larger number of operators than json. A most useful one is the “contains” operator @>. It works similar to the #> operator for json.

For example, let's find the entry that mentions Lucie, one of the mother-of-five's children:

```

test=# SELECT s.name,
        jsonb_pretty(sd.details_b) json
FROM student_details sd, students s
WHERE s.s_id = sd.s_id
AND sd.details_b @>
'{"merits":{"mother-of-five":{}}}' \gx

```

```

-[ RECORD 1 ]-----
name | Nina
json | {
      |   "merits": {
      |     "mother-of-five": {
      |       "Alex": "unknown",
      |       "Mark": "m",
      |       "Basil": "m",
      |       "Lucie": "f",
      |       "Simon": "m"
      |     }
      |   },
      |   "hobbies": "cosplay"
      | }

```

We have used the `jsonb_pretty()` function, which formats the output of the `jsonb` type.

Alternatively, you can use the `jsonb_each()` function, which expands key-value pairs:

```

test=# SELECT s.name,
         jsonb_each(sd.details_b)
FROM   student_details sd, students s
WHERE  s.s_id = sd.s_id
AND    sd.details_b @>
       '{"merits":{"mother-of-five":{}}}' \gx

```

```

-[ RECORD 1 ]-----
name      | Nina
jsonb_each | (hobbies,""cosplay"")
-[ RECORD 2 ]-----
name      | Nina
jsonb_each | (merits,{"mother-of-five":
          | {"Alex": "unknown", "Mark":
          | "m", "Basil": "m", "Lucie":
          | "f", "Simon": "m"}})

```

Note that the name of Nina's child is replaced by an empty space {} in the query. This syntax adds flexibility to the process of application development.

But what's more important, jsonb allows you to create indexes that support the @> operator, its inverse <@, and many other ones (the GIN index is typically the most efficient). The json type does not support indexes, so for high-load applications it is usually recommended to use jsonb.

To learn more about json and jsonb data types and their associated functions, see the PostgreSQL documentation at [postgrespro.com/doc/datatype-json](https://www.postgresql.org/docs/datatype-json) and [postgrespro.com/doc/functions-json](https://www.postgresql.org/docs/functions-json).

When the SQL:2016 standard was published, which included the SQL/JSON Path language, Postgres Professional developed its implementation, providing the jsonpath type. It was later committed to PostgreSQL 12.

- \$.a.b.c replaces the 'a'-'b'-'c' syntax that had to be used in PostgreSQL 11 or lower.
- The \$ symbol represents the current context element, that is, the JSON fragment that has to be parsed.
- @ represents the current context element in filter expressions. All the paths available in the \$ expression are searched.
- \* is a wildcard symbol. In expressions with \$ and @ it denotes any value of the path taking the hierarchy into account.
- \*\* is a wildcard that can denote any part of the path in expressions with \$ or @, without taking the hierarchy into account. It comes in handy if you don't know the exact nesting level of the elements.

- The ? operator is used to create a filter similar to WHERE. 133  
For example: \$.a.b.c ? (@.x > 10). ix

To find cosplay enthusiasts using the `jsonb_path_query()` function, you can write the following query:

```
test=# SELECT s_id, jsonb_path_query(  
    details::jsonb,  
    '$.hobbies ? (@ == "cosplay")'  
)  
FROM student_details;
```

```
 s_id | jsonb_path_query  
-----+-----  
 1556 | "cosplay"  
(1 row)
```

This query searches only through the JSON branch that begins with the “hobbies” key, checking whether the corresponding value equals “cosplay.” But if we replace “cosplay” with “guitarist,” the query won’t return anything because “guitarist” is used in our table as a key, not as a value of the nested element.

Queries can operate on two hierarchical levels: one inside the path expression, which defines the search area, and the other inside the filter expression, which matches the results with the specified condition. It means there are different ways to reach the same goal.

For example, the query

```
test=# SELECT s_id, jsonb_path_query(  
    details::jsonb,  
    '$.hobbies.guitarist.band?(@=="Postgressors")'  
)  
FROM student_details;
```

134 and the query

ix

```
test=# SELECT s_id, jsonb_path_query(
    details::jsonb,
    '$.hobbies.guitarist?(@.band=="Postgressors").band'
)
FROM student_details;
```

return the same result:

```
 s_id | jsonb_path_query
-----+-----
 1432 | "Postgressors"
(1 row)
```

In the first example, we defined a filter expression for each entry within the “hobbies.guitarist.band” branch. If we take a look at the JSON itself, we can see that this branch has only one value: “Postgressors.” So there was actually nothing to filter out. In the second example, the filter is applied one step higher, so we have to specify the path to the “group” within the filter expression; otherwise, the filter won’t find any values. If we use such syntax, we have to know the JSON hierarchy in advance. But what if we don’t know the hierarchy?

In this case, we can use the **\*\*** wildcard. It is an extremely useful feature! Suppose we are not sure what a “Strat” is: whether it’s a high-altitude balloon, a guitar, or a member of the highest social stratum. But we have to find out whether we have this word in our table at all. Previously, it would have required a complex search through the JSON document (unless we converted jsonb to text). Now you can simply run the following query:

```
test=# SELECT s_id, jsonb_path_exists(
    details::jsonb, '$.** ? (@ == "Strat")'
)
FROM student_details;
```

```
s_id | jsonb_path_exists
-----+-----
1451 | f
1432 | t
1556 | f
1451 | f
(4 rows)
```

You can learn more about SQL/JSON Path capabilities in the documentation ([postgrespro.com/doc/datatype-json#DATATYPE-JSONPATH](https://www.postgresql.org/docs/datatype-json#DATATYPE-JSONPATH)) and in the article “JSONPath in PostgreSQL: committing patches and selecting apartments” ([habr.com/en/company/postgrespro/blog/500440/](https://habr.com/en/company/postgrespro/blog/500440/)).

The SQL:2016 Standard defines a few more convenient features and expressions for working with JSON. Eventually, they make their way into PostgreSQL. We will discuss some of them here; for the full list, see the documentation: [postgrespro.com/doc/functions-json](https://www.postgresql.org/docs/functions-json).

Let’s get back to our friends Anna, Victor and Nina. They took part in a poll to determine what TV shows the younger generation likes. The results are stored in a table.

```
test=# CREATE TABLE tv_series(
    se_id int,
    s_id int REFERENCES students(s_id),
    se_details jsonb,
    CONSTRAINT pk_se PRIMARY KEY(s_id, se_id)
);

test=# INSERT INTO tv_series (se_id, s_id, se_details)
VALUES
(1, 1451,
'{"title": ["Game of Thrones",
            "Peaky Blinders"]}'),
```

```

136 (2, 1432,
ix   '{"title": "Babylon 5"}'
    ),
    (3, 1556,
     '{"title": ["Twin Peaks",
                  "Peaky Blinders"]}');

```

Database engineers analyze the results. They've got a powerful toolkit: aggregation functions, check constraints, an assortment of json(b) tools.

The arrays should first be split into separate elements (unnest). It can be done with jsonb\_array\_elements, but there's an issue:

```

text=# SELECT jsonb_array_elements(
        se_details->'title'
    )
FROM tv_series;

```

ERROR: cannot extract elements from a scalar

Naturally, the data from the questionnaires has been put into the database carelessly: Student 1432 does not have a corresponding array (of one element only, but still). Just to make sure, use IS [NOT] JSON ARRAY:

```

text=# SELECT se_id, se_details
FROM tv_series
WHERE se_details->'title' IS NOT JSON ARRAY \gx
-[ RECORD 1 ]-----
se_id      | 2
se_details | {"title": "Babylon 5"}

```

We can fix the type with jsonb\_set, which can modify jsonb structure:

```
text=# UPDATE tv_series
SET se_details = jsonb_set(
    se_details,
    '{title}',
    '["Babylon 5"]'
)
WHERE se_id = 2;
```

Now:

```
test=# SELECT DISTINCT jsonb_array_elements(
    se_details->'title'
)
FROM tv_series;

 jsonb_array_elements
-----
 "Twin Peaks"
 "Babylon 5"
 "Peaky Blinders"
 "Game of Thrones"
(4 rows)
```

Here is our final list. Note that IS JSON ARRAY belongs to a group of predicates, together with IS JSON SCALAR, IS JSON OBJECT and the more general IS JSON. See the documentation for more: [postgrespro.com/doc/functions-json](http://postgrespro.com/doc/functions-json).

## Integration with External Systems

Real-world applications are not isolated, and they often have to send data to each other. Such interactions can be implemented at the application level, for example, using web services or file exchange, or you can rely on the database functionality for this purpose.

138 PostgreSQL supports the ISO/IEC 9075-9 standard (SQL/MED,  
ix Management of External Data), which defines access to external data sources from SQL via a special mechanism of foreign data wrappers.

The idea is to access external (foreign) data as if it were located in regular PostgreSQL tables. It requires creating foreign tables, which do not contain any data themselves and only redirect all queries to an external data source. This approach facilitates application development, as it allows to abstract from specifics of a particular external source.

Creating a foreign table involves several steps.

1. The `CREATE FOREIGN DATA WRAPPER` command plugs in a library for working with a particular data source.
2. The `CREATE SERVER` command defines a foreign server. You should usually specify such connection parameters as host name, port number, and database name.
3. The `CREATE USER MAPPING` command provides username mapping since different PostgreSQL users can connect to one and the same foreign source on behalf of different external users.
4. The `CREATE FOREIGN TABLE` command creates foreign tables for the specified external tables and views, while `IMPORT FOREIGN SCHEMA` allows to import descriptions of some or all tables from the external schema.

Below we'll discuss PostgreSQL integration with the most popular databases: Oracle, MySQL, SQL Server, and PostgreSQL itself. But first we need to install the libraries required for working with these databases.

## Installing Extensions

The PostgreSQL distribution includes two foreign data wrappers: `postgres_fdw` and `file_fdw`. The first one is designed for working with external PostgreSQL databases, while the second one works with files on a server. Besides, the community develops and supports various libraries that provide access to many popular databases. To get the full list, take a look at [pgxn.org/tag/fdw](http://pgxn.org/tag/fdw).

Foreign data wrappers for Oracle, MySQL, and SQL Server are available as extensions:

- Oracle – [github.com/laurenz/oracle\\_fdw](https://github.com/laurenz/oracle_fdw)
- MySQL – [github.com/EnterpriseDB/mysql\\_fdw](https://github.com/EnterpriseDB/mysql_fdw)
- SQL Server – [github.com/tds-fdw/tds\\_fdw](https://github.com/tds-fdw/tds_fdw)

Follow the instructions on these web pages to build and install these extensions. If all goes well, you will see the corresponding foreign data wrappers in the list of available extensions. For example, for `oracle_fdw`:

```
test=# SELECT name, default_version
FROM pg_available_extensions
WHERE name = 'oracle_fdw' \gx
```

```
-[ RECORD 1 ]----+-----
name          | oracle_fdw
default_version | 1.2
```

### Oracle

First, let's create an extension, which will add a foreign data wrapper:

```
140 test=# CREATE EXTENSION oracle_fdw;
ix CREATE EXTENSION
```

Check that the corresponding wrapper has been added:

```
test=# \dew
```

```
List of foreign-data wrappers
-[ RECORD 1 ]-----
Name      | oracle_fdw
Owner     | postgres
Handler   | oracle_fdw_handler
Validator | oracle_fdw_validator
```

The next step is setting up a foreign server. In the `OPTIONS` clause, you have to specify the `dbserver` option, which defines connection parameters for the Oracle instance: server name, port number, and instance name.

```
test=# CREATE SERVER oracle_srv
      FOREIGN DATA WRAPPER oracle_fdw
      OPTIONS (dbserver '//localhost:1521/orcl');
CREATE SERVER
```

The PostgreSQL user `postgres` will be connecting to the Oracle instance as `scott`.

```
test=# CREATE USER MAPPING FOR postgres
      SERVER oracle_srv
      OPTIONS (user 'scott', password 'tiger');
CREATE USER MAPPING
```

We'll import foreign tables into a separate schema. Let's create it:

141  
ix

```
test=# CREATE SCHEMA oracle_hr;
CREATE SCHEMA
```

Now let's import some foreign tables. We'll do it for just two popular tables, dept and emp:

```
test=# IMPORT FOREIGN SCHEMA "SCOTT"
      LIMIT TO (dept, emp)
      FROM SERVER oracle_srv
      INTO oracle_hr;
IMPORT FOREIGN SCHEMA
```

Note that Oracle data dictionary stores object names in uppercase, while PostgreSQL system catalog saves them in lowercase. When working with external data in PostgreSQL, you have to double-quote uppercase Oracle schema names to avoid their conversion to lowercase.

Let's view the list of foreign tables:

```
test=# \det oracle_hr.*
      List of foreign tables
 Schema | Table | Server
-----+-----+-----
oracle_hr | dept | oracle_srv
oracle_hr | emp | oracle_srv
(2 rows)
```

Now run the following queries on the foreign tables to access the external data:

```
test=# SELECT * FROM oracle_hr.emp LIMIT 1 \gx
```

```

142  -[ RECORD 1 ]-----
ix    empno   | 7369
      ename  | SMITH
      job   | CLERK
      mgr   | 7902
      hiredate | 1980-12-17
      sal   | 800.00
      comm  |
      deptno | 20

```

Write operations on external data are also allowed:

```

test=# INSERT INTO oracle_hr.dept(deptno, dname, loc)
      VALUES (50, 'EDUCATION', 'MOSCOW');

```

```

INSERT 0 1

```

```

test=# SELECT * FROM oracle_hr.dept;

```

```

 deptno | dname      | loc
-----+-----+-----
      10 | ACCOUNTING | NEW YORK
      20 | RESEARCH   | DALLAS
      30 | SALES      | CHICAGO
      40 | OPERATIONS | BOSTON
      50 | EDUCATION  | MOSCOW

```

```

(5 rows)

```

## MySQL

Create an extension for the required foreign data wrapper:

```

test=# CREATE EXTENSION mysql_fdw;

```

```

CREATE EXTENSION

```

The foreign server for the external instance is defined by the host and port parameters:

```
test=# CREATE SERVER mysql_srv
      FOREIGN DATA WRAPPER mysql_fdw
      OPTIONS (host 'localhost', port '3306');
CREATE SERVER
```

We are going to establish connections on behalf of a MySQL superuser:

```
test=# CREATE USER MAPPING FOR postgres
      SERVER mysql_srv
      OPTIONS (username 'root', password 'p@ssw0rd');
CREATE USER MAPPING
```

The wrapper supports the IMPORT FOREIGN SCHEMA command, but we can also create a foreign table manually:

```
test=# CREATE FOREIGN TABLE employees (
      emp_no      int,
      birth_date  date,
      first_name  varchar(14),
      last_name   varchar(16),
      gender      varchar(1),
      hire_date   date)
      SERVER mysql_srv
      OPTIONS (dbname 'employees',
              table_name 'employees');
CREATE FOREIGN TABLE
```

Check the result:

```
test=# SELECT * FROM employees LIMIT 1 \gx
-[ RECORD 1 ]-----
emp_no      | 10001
birth_date  | 1953-09-02
first_name  | Georgi
last_name   | Facello
gender      | M
hire_date   | 1986-06-26
```

144 Just like the Oracle wrapper, `mysql_fdw` allows both read and  
ix write operations.

## SQL Server

Create an extension for the required foreign data wrapper:

```
test=# CREATE EXTENSION tds_fdw;  
CREATE EXTENSION
```

Create a foreign server:

```
test=# CREATE SERVER sqlserver_srv  
FOREIGN DATA WRAPPER tds_fdw  
OPTIONS (servername 'localhost', port '1433',  
database 'AdventureWorks');  
CREATE SERVER
```

The required connection information is the same: you have to provide the host name, the port number, and the database name. But the `OPTIONS` clause takes different parameters as compared to `oracle_fdw` and `mysql_fdw`.

We are going to establish connections on behalf of an SQL Server superuser:

```
test=# CREATE USER MAPPING FOR postgres  
SERVER sqlserver_srv  
OPTIONS (username 'sa', password 'p@ssw0rd');  
CREATE USER MAPPING
```

Let's create a separate schema for foreign tables:

```
test=# CREATE SCHEMA sqlserver_hr;  
CREATE SCHEMA
```

Import the whole HumanResources schema into the created PostgreSQL schema:

145  
ix

```
test=# IMPORT FOREIGN SCHEMA HumanResources
      FROM SERVER sqlserver_srv
      INTO sqlserver_hr;
IMPORT FOREIGN SCHEMA
```

You can display the list of imported tables using the `\det` command, or find them in the system catalog by running the following query:

```
test=# SELECT ft.ftrelid::regclass AS "Table"
      FROM pg_foreign_table ft;
```

Table

```
-----
sqlserver_hr.Department
sqlserver_hr.Employee
sqlserver_hr.EmployeeDepartmentHistory
sqlserver_hr.EmployeePayHistory
sqlserver_hr.JobCandidate
sqlserver_hr.Shift
(6 rows)
```

Object names are case-sensitive, so they should be enclosed in double quotes in PostgreSQL queries:

```
test=# SELECT "DepartmentID", "Name", "GroupName"
      FROM sqlserver_hr."Department"
      LIMIT 4;
```

DepartmentID	Name	GroupName
1	Engineering	Research and Development
2	Tool Design	Research and Development
3	Sales	Sales and Marketing
4	Marketing	Sales and Marketing

(4 rows)

146 Currently tds\_fdw supports only reading; write operations  
ix are not allowed.

## PostgreSQL

Create an extension and a wrapper:

```
test=# CREATE EXTENSION postgres_fdw;  
CREATE EXTENSION
```

We are going to connect to another database of the same server instance, so we only have to provide the dbname parameter when creating a foreign server. Other parameters (such as host, port, etc.) can be omitted.

```
test=# CREATE SERVER postgres_srv  
      FOREIGN DATA WRAPPER postgres_fdw  
      OPTIONS (dbname 'demo');  
CREATE SERVER
```

There is no need to specify the password if you create a user mapping within a single cluster:

```
test=# CREATE USER MAPPING FOR postgres  
      SERVER postgres_srv  
      OPTIONS (user 'postgres');  
CREATE USER MAPPING
```

Import all tables and views of the bookings schema:

```
test=# IMPORT FOREIGN SCHEMA bookings  
      FROM SERVER postgres_srv  
      INTO public;  
IMPORT FOREIGN SCHEMA
```

Check the result:

```
test=# SELECT * FROM bookings LIMIT 3;
```

book_ref	book_date	total_amount
000004	2015-10-12 14:40:00+03	55800.00
00000F	2016-09-02 02:12:00+03	265700.00
000010	2016-03-08 18:45:00+03	50900.00
000012	2017-07-14 09:02:00+03	37900.00
000026	2016-08-30 11:08:00+03	95600.00

(5 rows)

To learn more about `postgres_fdw`, see the documentation:  
[postgrespro.com/doc/postgres-fdw](https://postgrespro.com/doc/postgres-fdw).

Foreign data wrappers are also worth mentioning as the community considers them to be the basis for built-in sharding in PostgreSQL. Sharding is similar to partitioning: they both use a particular criterion to split a table into several parts that are stored independently. The difference is that partitions are stored on the same server, while shards are located on different ones. Partitioning has been available in PostgreSQL for quite a long time. Starting from version 10, this mechanism is being actively developed, and many useful features have already been added: declarative syntax, dynamic partition pruning, support for parallel operations, and other miscellaneous enhancements. You can also use foreign tables as partitions, which virtually turns partitioning into sharding.

However, significant work remains before sharding is fully practical:

- Consistency is not guaranteed: external data is accessed in separate local transactions rather than in a single distributed one.

- You can't duplicate the same data on different servers to enhance fault tolerance.
- All actions required to create tables on shards and the corresponding foreign tables have to be done manually.

These challenges are already addressed in Shardman, a new tool developed by Postgres Professional ([postgrespro.ru/products/shardman](http://postgrespro.ru/products/shardman)).

Another extension included into the distribution for working with PostgreSQL databases is dblink. It allows you to explicitly manage connections (to connect and disconnect), execute queries, and get the results asynchronously: [postgrespro.com/doc/dblink](http://postgrespro.com/doc/dblink).

# X Education and Certification

## Documentation

Reading the documentation is indispensable for professional use of PostgreSQL. It describes all the database features and provides an exhaustive reference that should always be readily available. Here, you can get full and precise information firsthand: it is written by developers themselves and is carefully kept up-to-date at all times. The PostgreSQL documentation is available at [www.postgresql.org/docs](http://www.postgresql.org/docs) or [postgrespro.com/docs](http://postgrespro.com/docs).

At Postgres Professional, we have translated the whole PostgreSQL documentation set into Russian, including the latest version. It is available on our website: [postgrespro.ru/docs](http://postgrespro.ru/docs).

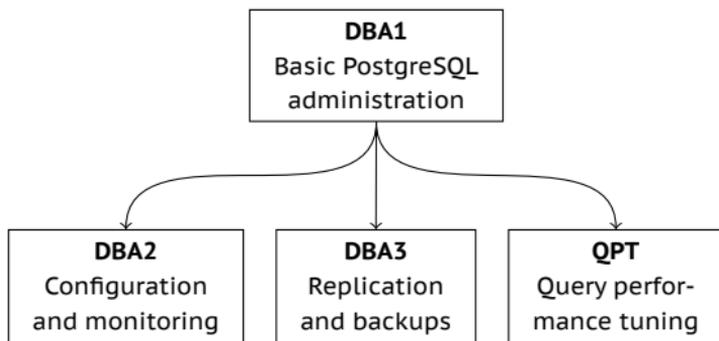
While working on this translation, we also compiled an English-Russian glossary, published at [postgrespro.ru/education/glossary](http://postgrespro.ru/education/glossary). We recommend consulting this glossary when translating English articles into Russian to use consistent terminology familiar to a wide audience.

There are also French ([docs.postgresql.fr](http://docs.postgresql.fr)), Japanese ([www.postgresql.jp/document](http://www.postgresql.jp/document)), and Chinese ([postgres.cn/docs](http://postgres.cn/docs)) translations provided by national communities.

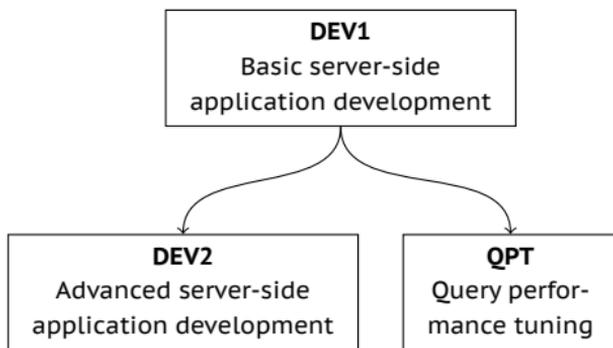
## Training Courses

We develop training courses for those who start using PostgreSQL or would like to improve their professional skills.

Courses for database administrators:

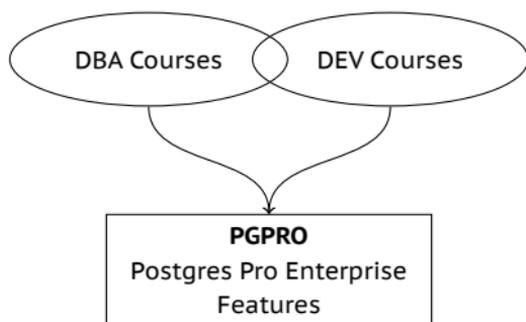


Courses for application developers:



An additional course about our primary product for those who have completed the administration and development courses:

151  
x



The PostgreSQL and Postgres Pro documentation contains every detail about PostgreSQL, but the information is often scattered across different chapters, so you may have to carefully read it several times before you gain full understanding.

Training courses are intended to complement the documentation rather than replace it. They consist of separate modules that gradually explain a particular topic, focusing on important practical information. Courses can broaden your outlook, structure the previously gained bits of knowledge, and help you find your way around the documentation, should you need to quickly get some particular details.

Each course topic includes theory and practice. In most cases, theory includes both slides and a live demo on a real system. Students receive all course slides with extensive annotations, outputs of demo scripts, solutions to practical assignments, and additional reference materials on select topics.

## Where and How to Take a Training

For non-commercial use and self-study, all course materials, including videos, are available on our website for free. You can find their Russian version at [postgrespro.ru/education/courses](https://postgrespro.ru/education/courses).

The courses currently translated into English are published at [postgrespro.com/community/courses](https://postgrespro.com/community/courses).

You can also take these courses at a specialized training center under the supervision of an experienced lecturer. Upon completing the course, you will receive a certificate of completion. Authorized training centers are listed here: [postgrespro.ru/education/where](https://postgrespro.ru/education/where).

### DBA1. Basic PostgreSQL administration

Duration: 3 days

Background knowledge required:

- Basic knowledge of databases and SQL.
- Familiarity with Unix.

Knowledge and skills gained:

- General understanding of PostgreSQL architecture.
- Installation, initial setup, server management.
- Logical structure and physical data layout.
- Basic administration tasks.
- User and access management.
- Backup, recovery, and replication.

### **Basic toolkit**

1. Installation and server management
2. Using psql
3. Configuration

### **Architecture**

4. PostgreSQL overview
5. Isolation and multi-version concurrency control
6. Vacuum
7. Buffer cache and write-ahead log

### **Data management**

8. Databases and schemas
9. System catalog
10. Tablespaces
11. Low-level

### **Administration tasks**

12. Monitoring

### **Access control**

13. Roles and attributes
14. Privileges
15. Row-level security
16. Connection and authentication

### **Backups**

17. Overview

### **Replication**

18. Overview

Course materials in English are available for self-study at [postgrespro.com/community/courses/DBA1](https://postgrespro.com/community/courses/DBA1).

## **DBA2. Configuring and monitoring PostgreSQL**

Duration: 4 days

Background knowledge required:

SQL fundamentals.

Good command of Unix OS.

Familiarity with PostgreSQL within the scope of the DBA1 course.

Knowledge and skills gained:

Setting up various configuration parameters based on the understanding of server internals.

Monitoring server activity and using the collected data for iterative tuning of PostgreSQL configuration.

Configuring localization settings.

Managing extensions and getting started with server upgrades.

Topics:

### **Multi-version concurrency control**

1. Transaction isolation
2. Pages and row versions
3. Data snapshots
4. HOT updates
5. Vacuum
6. Autovacuum item Freezing

### **Logging**

7. Buffer cache
8. Write-ahead log
9. Checkpoints
10. WAL configuration

**Locking**

11. Object locks
12. Row-level locks
13. Memory locks

**Administration tasks**

14. Managing extensions
15. Localization
16. Server upgrades

Course materials in Russian are available for self-study at [postgrespro.ru/education/courses/DBA2](https://postgrespro.ru/education/courses/DBA2).

**DBA3. Replication and backups**

Duration: 2 days

Background knowledge required:

SQL fundamentals.

Good command of Unix OS.

Familiarity with PostgreSQL within the scope of the DBA1 course.

Knowledge and skills gained:

Taking backups.

Setting up physical and logical replication.

Recognizing replication use cases.

Understanding cluster technologies.

Topics:

**Backups**

1. Logical backup
2. Base backup
3. WAL archive

**Replication**

4. Physical replication
5. Switchover to a replica
6. Logical replication
7. Usage scenarios

**Cluster Technologies**

8. Overview

Course materials in Russian are available for self-study at [postgrespro.ru/education/courses/DBA3](https://postgrespro.ru/education/courses/DBA3).

**DEV1. Basic server-side application development**

Duration: 4 days

Background knowledge required:

SQL fundamentals.

Experience with any procedural programming language.

Basic knowledge of Unix OS.

Knowledge and skills gained:

General information about PostgreSQL architecture.

Using the main database objects.

Programming in SQL and PL/pgSQL on the server side.

Using the main data types, including records and arrays.

Setting up client-server communication channels.

### **Basic toolkit**

1. Installation and server management, psql

### **Architecture**

2. A general overview of PostgreSQL
3. Isolation and MVCC
4. Buffer cache and WAL

### **Data organization**

5. Logical structure
6. Physical structure

### **Bookstore application**

7. Application schema and interface

### **SQL**

8. Functions
9. Procedures
10. Composite types

### **PL/pgSQL**

11. Overview and programming structures
12. Executing queries
13. Cursors
14. Dynamic commands
15. Arrays
16. Error handling
17. Triggers
18. Debugging

### **Access control**

19. Access control overview

### **Backup**

20. Logical backup

Course materials in English are available for self-study at [postgrespro.com/community/courses/DEV1](https://postgrespro.com/community/courses/DEV1).

## **DEV2. Advanced server-side application development**

Duration: 4 days

Background knowledge required:

General understanding of PostgreSQL architecture.  
Strong SQL and PL/pgSQL skills.  
Basic knowledge of Unix OS.

Knowledge and skills gained:

Understanding server internals.  
Using all PostgreSQL capabilities in application logic implementations.  
Extending database functionality to address specific tasks.

Topics:

### **Architecture**

1. Isolation
2. Server internals
3. Vacuum
4. Write-ahead logging
5. Locks

### **Bookstore**

6. Bookstore application 2.0

## **Extensibility**

159  
x

7. Connection pooling
8. Data types for large values
9. User-defined data types
10. Operator classes
11. Semi-structured data
12. Background processes
13. Asynchronous processing
14. Creating extensions
15. Programming languages
16. Aggregate and window functions
17. Full-text search
18. Physical replication
19. Logical replication
20. Foreign data

Course materials in Russian are available for self-study at [postgrespro.ru/education/courses/DEV2](https://postgrespro.ru/education/courses/DEV2).

## **QPT. Query Performance Tuning**

Duration: 2 days

Background knowledge required:

Familiarity with Unix OS.

Good command of SQL.

Some knowledge of PL/pgSQL will be useful, but is not mandatory.

Familiarity with PostgreSQL within the scope of the DBA1 course (for DBAs) or DEV1 (for developers).

160 Knowledge and skills gained:

x

In-depth understanding of query planning and execution.  
Performance tuning of the server instance.  
Troubleshooting query issues and optimizing queries.

Topics:

1. Airlines database
2. Query execution
3. Sequential scans
4. Index scans
5. Bitmap scans
6. Nested loop joins
7. Hash joins
8. Merge joins
9. Statistics
10. Query profiling
11. Optimization methods

Course materials in Russian are available for self-study at [postgrespro.ru/education/courses/QPT](https://postgrespro.ru/education/courses/QPT).

## **PGPRO. Postgres Pro Enterprise Features**

Duration: 2 days

Background knowledge required:

Familiarity with Unix OS.  
Good command of SQL.  
Some knowledge of PL/pgSQL will be useful, but is not mandatory.

Familiarity with PostgreSQL within the scope of the courses DBA1, DBA2 and QPT (for DBAs) or DEV1, DEV2 and QPT (for developers). 161  
x

Knowledge and skills gained:

Proficiency with Postgres Pro Enterprise signature features.

Topics:

1. Editions and features
2. Installation, configuration, upgrade
3. Transaction management
4. CFS – compressed file system
5. Query optimization
6. Adaptive optimization
7. Performance analysis
8. Load reporting – pgpro\_pwr
9. User profiles
10. Audit
11. Task planner
12. Backup – 1
13. Backup – 2
14. Backup – 3
15. Synchronized cluster – multimaster

Course materials in Russian are available for self-study at:  
[postgrespro.ru/education/courses/PGPRO](https://postgrespro.ru/education/courses/PGPRO).

## Professional Certification

The certification program, which was launched in 2019, is useful for both database professionals and their employers. Holding a certificate can give you an advantage when job hunting or negotiating your salary. Besides, it's a good opportunity to get an impartial evaluation of your knowledge.

For employers, the certification program facilitates recruiting, enables verification of PostgreSQL expertise of the current employees, and provides a means to control the quality of knowledge received in external employee training or to check the competence of third-party vendors and partners.

PostgreSQL certification is currently available only for database administrators, but in the future we plan to launch certification programs for PostgreSQL application developers too.

We offer three levels of certification, all of which require you to pass several tests.

**Professional level** confirms the knowledge in the following fields:

- General understanding of PostgreSQL architecture.
- Server installation, working in psql, tuning configuration settings.
- Logical and physical data structure.
- User and access management.
- General understanding of backup and replication concepts.

To obtain a certificate, you must successfully pass a test on the DBA1 course. For new applicants, the PostgreSQL 13 test version is recommended.

**Expert level** additionally confirms the knowledge in the following fields:

- PostgreSQL internals.
- Server setup and monitoring, database maintenance tasks.
- Performance optimization tasks, query tuning.
- Taking backups.
- Physical and logical replication setup for various usage scenarios.

There are two certification pathways:

- Have the Professional level for PostgreSQL 13 and pass DBA2-13, DBA3-13 and QPT-13 tests (in any order).
- Have the Expert level for PostgreSQL 10 and pass the transition test Expert 10–13.

**Master level** additionally confirms practical skills required for PostgreSQL database administration.

To get a certificate, it is required to have a certificate of the Expert level and successfully pass a hands-on test. This certification is currently under development.

Create an account under [postgrespro.ru/user](https://postgrespro.ru/user) and sign up for a certification test in your profile.

To pass a test, you should have:

- a good command of the corresponding courses and the documentation sections they refer to
- hands-on experience in working with PostgreSQL via psql

While taking the test, you can refer to our course materials and the PostgreSQL documentation, but usage of any other sources of information is prohibited.

- 164     Achieving a particular level is acknowledged by a certificate.  
x       Certificates do not expire, but they are tied to a specific PostgreSQL version and will become obsolete as that version is deprecated. So in several years you may want to take a test for a more recent PostgreSQL version.

To learn more about certification, visit [postgrespro.ru/education/cert](https://postgrespro.ru/education/cert).

## Academic Collaboration

Our company is committed to cultivating the next generation of database experts. This requires efforts at the earliest level of education, and it is only possible through collaboration with universities, colleges, and schools.

### Academic License

For educational organizations, we provide a free academic license for the Postgres Pro Standard DBMS. To apply, send a request to [academy@postgrespro.ru](mailto:academy@postgrespro.ru).

### Student Internships

We set up student labs at universities to provide internship opportunities. Our experts help the students select a graduation thesis related to PostgreSQL, develop and defend it.

### Competitions

Our company participates in organizing hackathons and various competitions for school and university students, including:

- the "Postgres Pro DBMS" section at the IT-Planet Competition: [challenge.braim.org/landing/postgres\\_contest](http://challenge.braim.org/landing/postgres_contest),
- the Russian National Open-Source Project Competition: [foss.kruzhok.org](http://foss.kruzhok.org).

Students who are winners or finalists of competitions hosted by the company or with our involvement, as well as lecturers from partner universities, can take professional certification tests for free. To receive a code for the free certification test, contact us at [certification@postgrespro.ru](mailto:certification@postgrespro.ru) and attach a supporting document.

## Academic Courses

We offer several academic courses produced in cooperation with professors from leading universities. All the courses can be used in educational institutions for free. Lecturers can use textbooks, slides, lecture videos, and other educational materials published on our website: [postgrespro.ru/education/university](http://postgrespro.ru/education/university).

Postgres Professional has contributed to several courses taught in such universities as Lomonosov Moscow State University, Higher School of Economics, Moscow Aviation Institute, Reshetnev Siberian State University of Science and Technology, and Siberian Federal University. Contact us if you are a university representative and would like to add database courses to the curriculum.

We also seek partnership with teachers and instructors who are ready to develop new original PostgreSQL courses. On our part, we provide all the required support and advice, edit the manuscripts and drive them to publication, as well as make arrangements for open lectures of the course authors in top Russian universities.

## SQL Basics

Course participants will learn about PostgreSQL and will be able to start working with it right away; no prior training is required. Starting with simple SQL queries, students will gradually get to more complex constructs, learn about transactions and query optimization.

This course is based on the following textbook (published in Russian):

E. Morgunov, **PostgreSQL. SQL Basics**. St. Petersburg : BHV-Petersburg, 2018.

ISBN 978-5-9775-4022-3



### Contents:

- Introduction
- Configuring the environment
- Basic operations
- Data types
- DDL fundamentals
- Queries
- Data manipulation
- Indexes
- Transactions
- Performance tuning

A soft copy of this book in Russian is available on our website:  
[postgrespro.ru/education/books/sqlprimer](http://postgrespro.ru/education/books/sqlprimer).

This course includes 36 hours of lectures and hands-on training. The author has been teaching it at top universities in Moscow and Krasnoyarsk for several years. You can download the course materials in Russian at [postgrespro.ru/education/university/sqlprimer](http://postgrespro.ru/education/university/sqlprimer).

**Evgeny Morgunov**, Ph.D in Technical Sciences, associate professor at the Informatics and Computer Science Department of Reshetnev Siberian State University of Science and Technology.



Evgeny lives in Krasnoyarsk. Before joining the university in 2000, he worked as a programmer for over 10 years, including developing a banking application system. He got to learn PostgreSQL in 1998. Being an advocate of using free open-source software in academic activities, he has initiated the use of PostgreSQL and FreeBSD operating system as part of the “Programming Technology” course. Evgeny is a member of the International Society for Engineering Pedagogy (IGIP). He has been using PostgreSQL in teaching for more than 20 years.

## Database Technology Fundamentals

A modern academic course that combines in-depth theory with relevant practical skills of database design and deployment.



The first part contains the key information about database management systems: relational data model, the SQL language, transaction processing.

The second part dives into the underlying database technology and its development trends. Some topics covered in the first part are discussed again at a deeper level.

Contents:

### **Part I. From Theory to Practice**

- Introduction
- Some database theory
- Getting started with databases
- Introduction to SQL
- Database access management
- Transactions and data consistency
- Database application development
- Relational model extensions
- Various types of database systems

## Part II. From Practice to Proficiency

- Database system architecture
- Storage structures and the main algorithms
- Query execution and optimization
- Transaction management
- Database reliability
- Advanced SQL features
- Database functions and procedures
- PostgreSQL extensibility
- Full-text search
- Data security
- Database administration
- Replication
- Parallel and distributed database systems

A soft copy of this book in Russian is available on our website: [postgrespro.ru/education/books/dbtech](http://postgrespro.ru/education/books/dbtech).

This course offers 24 hours of lectures and 8 hours of hands-on training. It was delivered by Boris Novikov at the faculty of Computational Mathematics and Cybernetics of Lomonosov Moscow State University. You can download the course materials in Russian at [postgrespro.ru/education/university/dbtech](http://postgrespro.ru/education/university/dbtech).

**Boris Novikov**, Dr. Sci. in Physics and Mathematics, professor at the Informatics Department of Higher School of Economics in St. Petersburg.

His academic interests mainly concern various aspects of designing, developing, and deploy-



170 ing database systems and applications, as well as scalable  
x distributed systems for Big Data processing and analytics.

**Ekaterina Gorshkova**, Ph.D. in Physics and Mathematics.

An expert in designing high-load data-intensive applications. Her academic interests include machine learning, data-flow analysis, and data retrieval.

**Natalia Grafeeva**, Ph.D. in Physics and Mathematics, associate professor at the Informatics and Data Analysis Department of St. Petersburg State University.

Her academic interests include databases, data retrieval, Big Data, and smart data analysis. She is an expert in information system design, development, and maintenance, as well as in course design and teaching.

## Books

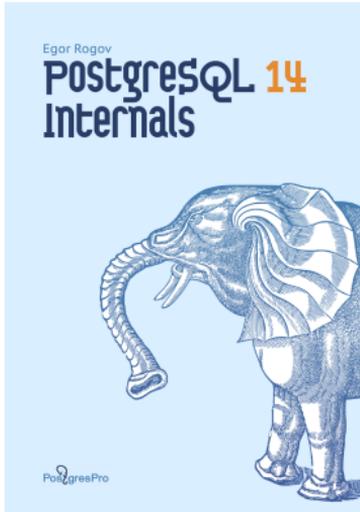
### PostgreSQL Internals

This book is for those who will not settle for a black-box approach when working with a database. Targeted at readers who have some experience with PostgreSQL, this book will also be useful for those who are familiar with another database system but switch over to PostgreSQL and would like to understand how they differ.

E. Rogov, **PostgreSQL 14 Internals**. Moscow : DMK Press, 2022

ISBN 978-5-6045970-4-0 (in English)

ISBN 978-5-93700-305-8 (in Russian)



You will not find any ready-made recipes in this book. But the provided explanations of the inner mechanics will enable you to critically evaluate other people's experience and come to your own conclusions. The author goes into details of PostgreSQL internals and shows how to run experiments and verify information that inevitably gets deprecated.

171  
x

**Egor Rogov** has been working in the education department in Postgres Professional since 2015; he develops and teaches training courses, publishes blog posts, writes and edits books.



Contents:

Introduction

### **Part I. Isolation and MVCC**

Isolation • Pages and Tuples • Snapshots • Page pruning and HOT updates • Vacuum and autovacuum • Freezing • Rebuilding tables and indexes

### **Part II. Buffer cache and WAL**

Buffer cache • Write-ahead log • WAL modes

### **Part III. Locks**

Relation-level locks • Row-level locks • Miscellaneous locks • Locks on memory structures

### **Part IV. Query execution**

Query execution stages • Statistics • Table access methods • Index access methods • Index scans • Nested loop • Hashing • Sorting and Merging

### **Part V. Index types**

Hash • B-tree • GiST • SP-GiST • GIN • BRIN

A soft copy of this book is available on our website:  
[postgrespro.com/community/books/internals](https://postgrespro.com/community/books/internals).

## **PostgreSQL Monitoring**

Monitoring is a critical aspect of PostgreSQL administration. The book encompasses this topic, providing vast information on available tools, tips on how to use them and ways of processing the information you get from them.

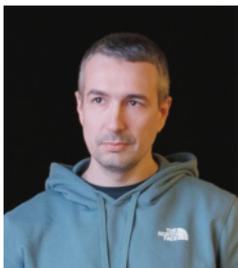
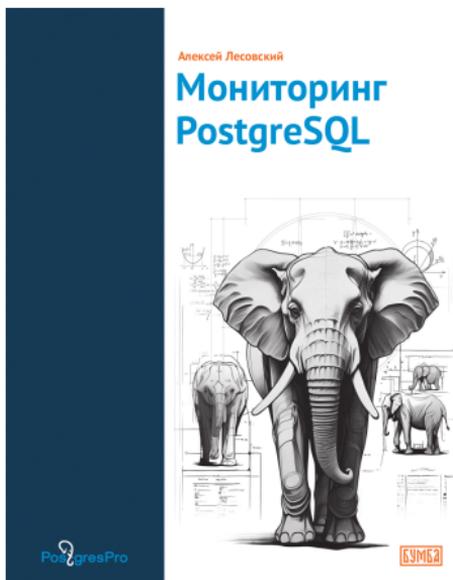
This book provides insights into PostgreSQL operations and monitoring techniques, helping readers optimize database performance and effectively address administrative challenges.

A. Lesovsky, **PostgreSQL Monitoring**. Moscow : Bumba, 2024

ISBN 978-5-907754-42-3

## Contents:

- Introduction
- About this book
- Statistics overview
- Activity statistics
- Query and function execution
- Databases
- Shared memory and input/output
- Write-ahead Log
- Replication
- Vacuum
- Operation execution process
- Appendix: Testing environment



**Alexey Lesovsky** is an expert database administrator, system administrator, developer and devops engineer with almost 20 years of experience in managing large-scale complex systems as well as in software design and development.

The book comes together with a docker environment for you to test the examples and run experiments in. It is an

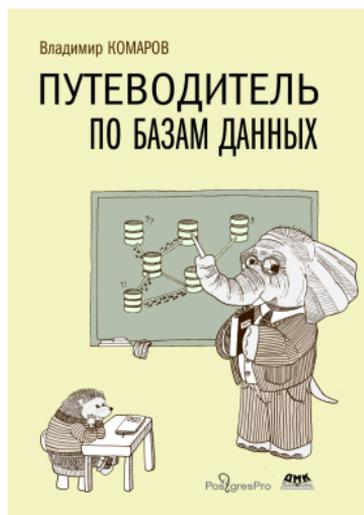
174 interesting read for database administrators, system administrators, reliability experts and performance optimization enthusiasts.

X

A soft copy of this book in Russian is available on our website: [postgrespro.ru/education/books/monitoring](http://postgrespro.ru/education/books/monitoring).

## A Database Guide

The book outlines the architectural foundations of all modern database management systems and explains the algorithms and data structures they utilize. Special attention is paid to the implementation of the same principles in functionally similar platforms.



A must-read for anyone dissatisfied with their “Software Engineering in Three Months” course. It supplements purely practical skills by giving understanding of underlying patterns. This is a book for software architects and senior developers—the elite and those aspiring to join their ranks.

V. Komarov, **A Database Guide**. Moscow : DMK Press, 2024

ISBN 978-5-93700-287-7

A soft copy of this book in Russian is available on our website:  
[postgrespro.ru/education/books/dbguide](http://postgrespro.ru/education/books/dbguide).

175  
x

**Vladimir Komarov** is an IT generalist: software developer, database administrator, data and infrastructure architect, lecturer and a bit of an evangelist.



Contents:

### **Part I. Database Classification**

Data models • Extra classification criteria

### **Part II. Data Access**

Storage structures • Data processing

### **Part III. DBMS Architecture**

Data consistency guarantees • DBMS structure

### **Part IV. Distributed Databases**

The distribution compromise

Modifying distributed data

### **Part V. Recovery**

Replication • Backup

### **Part VI. Database Operation**

Database management • Hardware • Financials

### **Part VII. Database Security**

Access control • Internal threats



# XI The Hacker's Guide to the Galaxy

## News and Discussions

Anyone can follow PostgreSQL news, learn about the features planned for the next release, and stay up-to-date with current events.

Plenty of interesting and useful content is published in various related blogs. For example, the [planet.postgresql.org](http://planet.postgresql.org) website aggregates all the English-language articles in one place. Many articles in Russian can be found at [habr.com/hub/postgresql](http://habr.com/hub/postgresql), including those published by Postgres Professional. For some of our articles, an English translation is available at [habr.com/en/company/postgrespro/blog/](http://habr.com/en/company/postgrespro/blog/). There are also dedicated YouTube channels, such as [youtube.com/RuPostgres](http://youtube.com/RuPostgres) and [youtube.com/PostgresTV](http://youtube.com/PostgresTV).

There is also a Wiki project ([wiki.postgresql.org](http://wiki.postgresql.org)), where you can find FAQs, training materials, articles about system setup and optimization, migration specifics from different database systems, and much more.

Almost 14 000 Russian-speaking PostgreSQL users subscribe to the “pgsql–PostgreSQL” channel on Telegram ([t.me/pgsql](https://t.me/pgsql)), establishing an active and helpful community.

178 You can also ask your questions on [stackoverflow.com](https://stackoverflow.com). Do  
xi not forget to add the “postgresl” tag.

As for Postgres Professional news, they are published in its corporate blog at [postgrespro.com/blog](https://postgrespro.com/blog).

## Mailing Lists

To get all the news firsthand, without waiting for someone to write a blog post, you can subscribe to mailing lists. Following the tradition, PostgreSQL developers discuss all questions exclusively via email.

You can find all the mailing lists at [postgresql.org/list](https://postgresql.org/list). Some of them are:

- [pgsql-hackers](#) (typically called simply “hackers”), the main list for everything related to development
- [pgsql-general](#) used to discuss general questions
- [pgsql-bugs](#) for bug reports
- [pgsql-docs](#) for documentation
- [pgsql-translators](#) for translation-related discussions
- [pgsql-announce](#) to get new release announcements

and many more.

Once you sign up, you will start receiving regular emails and will be able to participate in discussions if you wish. Another option is to browse through the email archive at [postgresql.org/list](https://postgresql.org/list) or on our company’s website ([postgrespro.com/list](https://postgrespro.com/list)).

Another way to keep up with the news without spending too much time is to check the [commitfest.postgresql.org](https://commitfest.postgresql.org) page. Here the community opens the so-called commitfests for developers to submit their patches. For example, commitfest 01.03.2024–31.03.2024 was open for version 17, while commitfest 01.07.2024–31.07.2024 was related to the next version already. It allows the community to stop accepting new features at least about six months before the release and have the time to stabilize the code.

Patches undergo several stages: first, they are reviewed and fixed, and then they are either accepted, moved to the next commitfest, or rejected (if you are completely out of luck).

This way, you can stay informed about new features already included in PostgreSQL or planned for the next release.

## Conferences

Moscow, St. Petersburg, and other cities in Russia host the annual international conference **PGConf** ([pgconf.ru](https://pgconf.ru)), attended by hundreds of PostgreSQL users and developers.

Besides, several Russian cities host conferences on broader topics, including databases in general and PostgreSQL in particular. We will name only a few:

**HighLoad++** in Moscow and other cities ([highload.ru](https://highload.ru));

**CodeFest** in Novosibirsk ([codefest.ru](https://codefest.ru));

**DUMP** in St. Petersburg ([dump-spb.ru](https://dump-spb.ru)) and Ekaterinburg ([dump-ekb.ru](https://dump-ekb.ru));

**Stachka** in Ulyanovsk ([nastachku.ru](http://nastachku.ru));

**Heisenbug** in Moscow ([heisenbug.ru](http://heisenbug.ru));

**Sysconf** in Moscow ([sysconf.pro](http://sysconf.pro)).

Naturally, PostgreSQL conferences are held all over the world. The major ones are:

**PGCon** in Ottawa, Canada ([pgcon.org](http://pgcon.org))

**PGConf Europe** ([pgconf.eu](http://pgconf.eu))

The list of upcoming events can be found at [postgresql.org/about/events](http://postgresql.org/about/events).

In addition to conferences, there are less official regular meetups, including online ones.

# XII Postgres Professional

Postgres Professional company was founded in 2015; it unites key Russian developers whose contributions to PostgreSQL are recognized in the global community. Fostering database development expertise in Russia, the company currently employs about 400 developers, architects, engineers, and other specialists.

Postgres Professional company delivers several versions of Postgres Pro database system based on PostgreSQL, as well as develops new core features and extensions and provides support for application system design, maintenance, and migration to PostgreSQL.

The company pays much attention to education. It hosts PgConf.Russia, the largest international annual PostgreSQL conference in Moscow, and participates in other conferences all over the world.

Contact information:

7A Dmitry Ulyanov str., Moscow, Russia, 117036

+7 495 150-06-91

[info@postgrespro.ru](mailto:info@postgrespro.ru)

## Postgres Pro Database System

Postgres Pro is a Russian commercial database system developed by the Postgres Professional company. Based on the open-source PostgreSQL database system, Postgres Pro offers many additional features to satisfy the needs of enterprise customers. It is included in the unified register of Russian software.

**Postgres Pro Standard** includes all PostgreSQL features, along with additional extensions and core patches, including those that have not yet been accepted by the community. As a result, its users can get access to useful functionality and improve performance without having to wait for the next PostgreSQL version to be released.

**Postgres Pro Enterprise** is a considerably reworked version of the database system; offering better stability and increased performance, it can address challenging production-level tasks.

Both Postgres Pro versions have been extended with the required information security functionality and are **certified by FSTEC** (Federal Service for Technical and Export Control).

To use any Postgres Pro version, you must purchase a license. A trial version is available for free, and you can also obtain Postgres Pro at no cost for educational purposes or application development.

Educational institutions are eligible for a free academic Postgres Pro Standard license.

To learn more about the features specific to different Postgres Pro versions, go to [postgrespro.com/products](https://postgrespro.com/products).

## Postgres Pro Enterprise Manager

183  
xii

Postgres Pro Enterprise Manager (PPEM) is an integrated control panel for Postgres Pro Enterprise.

PPEM is an all-in-one console with a streamlined monitoring and control interface. It lets you perform general administration tasks in the browser and aggregates access to all your databases and instances. PPEM increases the database administration team's efficiency and speeds up daily administrative tasks.

For details, see: [postgrespro.ru/products/PPEM](https://postgrespro.ru/products/PPEM).

## Postgres Pro Backup Enterprise

Postgres Pro Backup Enterprise (pg\_probackup) is a cluster backup and recovery tool for PostgreSQL.

It utilizes a centralized backup catalog, stored locally or remotely (including an S3 storage).

Pro Backup supports page-level incremental backup, backup storage policies, integrity verification, point-in-time recovery, parallel backup and recovery, data compression, and CFS file systems.

For details, see: [postgrespro.ru/products/pg\\_probackup](https://postgrespro.ru/products/pg_probackup).

## Shardman

Shardman is a PostgreSQL-based distributed relational DBMS with all the extended features of Postgres Pro.

184      Shardman is PostgreSQL-compatible, provides strong guar-  
xii      antees of data isolation and consistency, enables horizontal  
scaling, and boasts built-in redundancy mechanisms.

For details, see: [postgrespro.ru/products/shardman](https://postgrespro.ru/products/shardman).

## Careers

We are always looking for people who want to help grow the Postgres ecosystem and the open source community at large.

The Postgres Professional team comprises the majority of Russian PostgreSQL contributors. We develop kernel-level code for PostgreSQL and its extensions, create ecosystem products for database monitoring and administration. Our experts develop DBaaS solutions, carry out technical audits and migrations, offer technical support, participate in research, and drive innovation.

We are looking for system software engineers, backend and frontend developers, DevOps, QA engineers, database administrators, and other specialists. For details, check out our career portal: [career.postgrespro.ru](https://career.postgrespro.ru).

## Services

### Fault-Tolerant Solutions for Postgres

Designing and implementing high-load, high-performance, and fault-tolerant production systems; providing consulting services. Deploying Postgres and optimizing system configuration.

## **Vendor Technical Support**

185  
xii

24x7 support for Postgres Pro and PostgreSQL: system monitoring, disaster recovery, incident analysis, performance management, debugging both core features and extensions.

## **Migration of Application Systems**

Estimating the complexity of migration to Postgres from other database systems. Defining the architecture and the required changes for new solutions. Migrating application systems to Postgres and providing support during migration.

## **Postgres Training**

Courses for database administrators, system architects, and application developers covering Postgres specifics and efficient use of its advantages.

## **Database System Audit**

Database system evaluation by Postgres Professional experts. Information security audit for Postgres-based systems.

A complete list of services is available at [postgrespro.com/services](https://postgrespro.com/services).

Pavel Luzanov  
Egor Rogov  
Igor Levshin

## **Postgres. The First Experience**

Translated by Liudmila Mantrova  
and Alexander Meleshko  
Cover design by Eteri Bezhashvili

11th edition, revised and updated

[postgrespro.com/community/books/introbook](https://postgrespro.com/community/books/introbook)

© Postgres Professional, 2016–2025

ISBN 978-5-6045970-8-8