

PostgreSQL Reference Card

Connection arguments

- `-h` : hostname or IP address or socket directory
- `-p` : port number
- `-U` : role name
- `-d` : database name
- connection string : `"host=<hostname> port=<number> user=<role> dbname=<name>"`
- service file : `~/.pg_service.conf`
- global service file : `pg_service.conf` in directory :
 - Redhat : `/etc/sysconfig/pgsql/`
 - Debian : `/etc/postgresql-common/`

Client psql

Arguments

- `-c '<sql>'` : run SQL or internal command
- `-f <file>` : run a file
- `-1` : run orders in a transaction
- `-s` : run orders step by step

Internal commands

- `\c [dbname] [username] [host] [port]`
- `\c connstring` : open a new connection
- `\conninfo` : display current connection
- `\l` : show databases
- `\du, \dg, \drg` : show roles and membership
- `\dn` : show namespaces (schemas)
- `\dt` : show visible tables (using `search_path`)
- `\dP` : show partitioned tables
- `\dx` : show installed extensions
- `\df` : show visible functions
- `\d[+] <object>` : describe [with details] object
- `\i <file>` : run a file
- `\timing [on|off]` : toggle display running time
- `\gx` : expand result's display
- `\gexec` : execute query's result
- `\watch` : execute in loop previous query
- `\?` : help about internal command
- `\h` : help about SQL orders

Directory and files

Creation with initdb

- `-D` : data directory (`$PGDATA`)
- `-k` : enable checksum in data pages
- `-X` : WAL directory

Upgrade with pg_upgrade

- `-b, -B` : old and new executable directory
- `-d, -D` : old and new data directory
- `-c` : check only
- `-k` : use hard link instead of copy

Default directories

- Redhat : `/var/lib/pgsql/18/data`
- Debian : `/var/lib/postgresql/18/main`

Contents

- `base` : database directories and data files
- `pg_wal` : WAL directory

Configuration

- `pg_hba.conf` : authentication file
- `postgresql.conf` : main configuration file
- `postgresql.auto.conf` : used by ALTER SYSTEM

Default directories

- Redhat : `/var/lib/pgsql/18/data`
- Debian : `/etc/postgresql/18/main`

Main parameters

- `shared_buffers` : main memory, shared between backends
- `work_mem` : memory used by sort or hash
- `maintenance_work_mem` : memory used by VACUUM
- `autovacuum_[vacuum|analyze]_scale_factor` : table ratio to trigger VACUUM or ANALYZE
- `max_wal_size, checkpoint_timeout` : trigger a CHECKPOINT
- `default_statistics_target` : statistics size
- `io_method` : allows the use of an input/output method (e.g. `uring`)

Example commands

- `SET work_mem = '100MB';`
- `ALTER SYSTEM SET work_mem = '100MB';`
- `SET search_path TO public, <namespace1>, <namespace2>;`
- `ALTER DATABASE <dbname> SET search_path FROM CURRENT;`
- `ALTER TABLE <table> SET (autovacuum_vacuum_scale_factor = 0.1);`
- `CREATE INDEX on <table> USING btree|gist|gin|brin (<column>, ... ) [WHERE <predicate>];`

Activity Log Directory

- Redhat : `/var/lib/pgsql/18/data/log`
- Debian : `/var/log/postgresql/`

Catalog

- `pg_settings` : configuration parameters
- `pg_stat_activity` : live sessions
- `pg_locks` : live locks
- `pg_stat_user_tables` : tables and statistics
- `pg_stat_replication` : replication status

Queries

Live queries

```
SELECT datname, pid, query_id, query_start,
       wait_event_type, wait_event,
       state, query
FROM pg_stat_activity
WHERE backend_type='client backend'
AND pid <> pg_backend_pid();
```

Live locks

```
SELECT datname, pid, backend_start, wait_event,
       state, relation::regclass, mode, granted
FROM pg_locks l
JOIN pg_stat_activity a USING (pid)
WHERE pid <> pg_backend_pid();
```

SQL functions

- `pg_reload_conf()` : reload configuration files
- `pg_cancel_backend(pid)` : cancel a query
- `pg_terminate_backend(pid)` : terminate a session
- `pg_create_restore_point(text)` : create a restore point in WAL for PITR
- `pg_is_in_recovery()` : is cluster a standby?
- `pg_wal_replay_pause()` : pause the streaming replication
- `pg_wal_replay_resume()` : resume the streaming replication
- `pg_create_physical_replication_slot(name)` : create a replication slot
- `pg_drop_replication_slot(name)` : drop a replication slot

Backups

Logical backup : `pg_dump`

- `-F p|c|t|d` :
 - plain : plain SQL text
 - custom : binary compressed file
 - tar : TAR file
 - directory : compressed directory
- `-j N` : jobs number
- `-f` : backup file
- `-c` : add object's DROP commands
- `-C` : add database creation
- `-n, -N` : include or exclude namespace
- `-t, -T` : include or exclude table
- `-s` : backup schema, not data
- `-a` : backup data, not schema
- `-Z` : gzip, lz4 or zstd compression

```
pg_dump -Fc -Z zstd -f dbname.20250925.dmp
pg_dumpall -g -f globals.sql
```

Logical restore : `pg_restore`

- `-F c|t|d` : backup format
- `-l` : catalog extraction
- `-L <file>` : use external catalog
- `-C` : create the database
- `-j N` : restore jobs number

```
pg_restore -l dbname.20250925.dmp \
> catalog.txt
pg_restore -L catalog.txt -j4 \
-C -d postgres dbname.20250925.dmp
```

Physical backup : `pg_basebackup`

- `-D <dest>` : backup destination
- `-F t|p` : backup format
- `-X fetch|stream|none` : WAL fetch method
- `-c fast|spread` : CHECKPOINT's type
- `-l text` : backup label
- `-R` : make standby.signal and write recovery configuration in postgresql.auto.conf
- `-P` : progress bar

```
pg_basebackup -Fp -X stream -c fast \
-l '20250925.backup' -R -P \
-D backup.20250925.d
```

Instance control from system

SystemD

```
systemctl restart|reload|start|stop|status \
postgresql-18
```

Debian

```
pg_ctlcluster 18 main \
restart|reload|start|stop|status
```


Debian scripts

- `pg_lsclusters` : show clusters
- `pg_createcluster` : create a new cluster
- `pg_dropcluster` : drop a cluster
- `pg_upgradecluster` : upgrade to a new major version

download me : <https://www.loxodata.com/refcard>



PostgreSQL Reference Card



contact@loxodata.com +33 1 797 2 5775
www.loxodata.com