

The config file(s)

Main config file is **postgresql.conf** in instance directory (or specified explicitly as startup argument).

Settings changes (incl. ALTER SYSTEM) are often not applied immediately: see the context property for how/when setting changes are applied.

Official doc: [19.1. Setting Parameters](#) / [52.24. pg_settings](#)

Connection/network

listen_addresses	interfaces to listen on (ipv4 and v6)
port	TCP port
max_connections	Max simultaneous sessions
ssl	Enable encrypted transport

Logging

log_destination	Log format
log_filename	Log filename pattern
log_min_messages	Minimum event severity to be logged
log_min_duration_statement	Log all statements with exec time > value
log_checkpoint	Log checkpoints
log_line_prefix	Prefix for all log lines
log_lock_waits	Log lock waits longer than value
log_statement	Log SQL statements (DML only, DDL only, ...)
log_temp_files	Log temp files larger than value

See [19.8. Error Reporting and Logging](#) for filename format and templates

WAL/checkpoints

wal_level	Transaction log level: minimal (no log), replica, logical
fsync	Force flushing OS disk cache to physical storage
wal_sync_method	Method used for forcing WAL updates out to disk when fsync is on
synchronous_commit	Wait for writes to be committed to disk before sending commit to client
full_page_writes	Do not write partial pages to WAL
wal_buffers	Size of WAL buffers in RAM

WAL/checkpoints (cont)

checkpoint_timeout	Max delay between checkpoints
checkpoint_completion_target	Spread checkpoint over n times checkpoint_timeout
min/max_wal_size	Soft limits (targets) for WAL size

Replication

archive_mode	Enable WAL archiving
archive_command	Command used to archive WAL files
max_wal_senders	Max concurrent replication threads
wal_keep_size	Size of preserved WAL segments on the primary
max_replication_slots	Number of replication slots (max concurrent standby connections)
hot_standby	When in standby (in recovery), accept connections
hot_standby_feedback	When in standby, give status feedback to primary

See [19.5.1. Write Ahead Log settings](#), [25.3.1. Setting Up WAL Archiving](#) and [19.6. Replication](#)

Resources

shared_buffers	Instance-wide shared memory (see oracle sga)
temp_buffers	Dedicated memory per session (see oracle pga)
work_mem	Dedicated memory per query
maintenance_work_mem	Memory reserved for maintenance ops (vacuum, index creation, etc)
autovacuum_work_mem	Memory dedicated to autovacuum processes



Autovacuum controls

autovacuum	Enable/disable autovacuum
autovacuum_max_workers	Max parallel autovac jobs
autovacuum_naptime	Max delay before autovac
autovacuum_vacuum_cost_limit	Accumulated I/O cost of autovac before pausing
autovacuum_vacuum_cost_delay	Duration of autovacuum pause (see previous)
autovacuum_vacuum_threshold	Insert/update threshold to trigger autovac (absolute)
autovacuum_vacuum_scale_factor	Insert/update threshold to trigger autovac (relative)

See [19.10. Automatic Vacuuming](#)

Autovacuum keeps count of I/O (as "cost"). Once cost has reached `autovacuum_vacuum_cost_limit`, it pauses for `autovacuum_vacuum_cost_delay`.

Vacuum threshold is the sum of `vacuum_threshold` and `vacuum_scale_factor`.

Autovacuum trigger thresholds have separate, identical parameters for inserts.

Planner/optimizer

seq_page_cost	Sequential page read cost
random_page_cost	Random page read cost
cpu_operator_cost	Operator call cost
effective_cache_size	Estimated system cache size
default_statistics_target	"Target quality" of the optimizer: higher values takes longer to calculate plan but yields better results
from_collaps_limit	Merge nested FROM (to avoid "reorder experimentation" from the planner) if sub-FROM would yield less than this setting
join_collaps_limit	See <code>from_collapse_limit</code> but with JOINS instead of FROMs

See [19.7. Query Planning](#)

