



# Что нового в PostgreSQL в плане мониторинга



PostgreSQL DBA, системный администратор

Data Egret — PostgreSQL консалтинг и поддержка


Люблю мониторинги, статистику, визуализации





**Nagios®**





**PgCluu**

**PgHero**

**PGObserver**



PGWATCH



Grafana



## **The Statistics Collector**

<https://www.postgresql.org/docs/current/monitoring-stats.html>



PostgreSQL's *statistics collector* is a subsystem that supports collection and reporting of information about server activity. Presently, the collector can count accesses to tables and indexes in both disk-block and individual-row terms. It also tracks the total number of rows in each table, and information about vacuum and analyze actions for each table. It can also count calls to user-defined functions and the total time spent in each one.

PostgreSQL also supports reporting dynamic information about exactly what is going on in the system right now, such as the exact command currently being executed by other server processes, and which other connections exist in the system. This facility is independent of the collector process.

### 27.2.1. Statistics Collection Configuration

Since collection of statistics adds some overhead to query execution, the system can be configured to collect or not collect information. This is controlled by configuration parameters that are normally set in `postgresql.conf`. (See [Chapter 19](#) for details about setting configuration parameters.)

The parameter `track_activities` enables monitoring of the current command being executed by any server process.

The parameter `track_counts` controls whether statistics are collected about table and index accesses.

The parameter `track_functions` enables tracking of usage of user-defined functions.

The parameter `track_io_timing` enables monitoring of block read and write times.

Normally these parameters are set in `postgresql.conf` so that they apply to all server processes, but it is possible to turn them on or off in individual sessions using the `SET` command. (To prevent ordinary users from hiding their activity from the administrator, only superusers are allowed to change these parameters with `SET`.)

The statistics collector transmits the collected information to other PostgreSQL processes through temporary files. These files are stored in the directory named by the `stats_temp_directory` parameter, `pg_stat_tmp` by default. For better performance, `stats_temp_directory` can be pointed at a RAM-based file system, decreasing physical I/O requirements. When the server shuts down cleanly, a permanent copy of the statistics data is stored in the `pg_stat` subdirectory, so that statistics can be retained across server restarts. When recovery is performed at server start (e.g., after immediate shutdown, server crash, and point-in-time recovery), all statistics counters are reset.

### 27.2.2. Viewing Statistics

Several predefined views, listed in [Table 27.1](#), are available to show the current state of the system. There are also several other views, listed in [Table 27.2](#), available to show the results of statistics collection. Alternatively, one can build custom views using the underlying statistics functions, as discussed in [Section 27.2.20](#).

When using the statistics to monitor collected data, it is important to realize that the information does not update instantaneously. Each individual server process transmits new statistical counts to the collector just before going idle; so a query or transaction still in progress does not affect the displayed totals. Also, the collector itself emits a new report at most once per `PGSTAT_STAT_INTERVAL` milliseconds (500 ms unless altered while building the server). So the displayed information lags behind actual activity. However, current-query information collected by `track_activities` is always up-to-date.

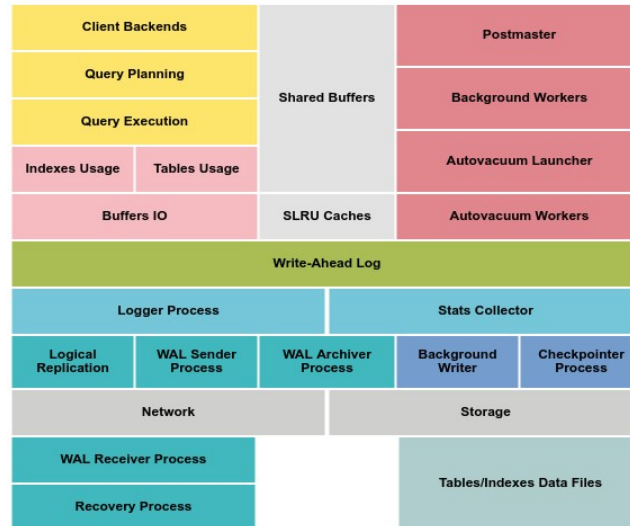
Another important point is that when a server process is asked to display any of these statistics, it first fetches the most recent report emitted by the collector process and then continues to use this snapshot for all statistical views and functions until the end of its current transaction. So the statistics will show static information as long as you continue the current transaction. Similarly, information about the current queries of all sessions is collected when any such information is first requested within a transaction, and the same information will be displayed throughout the transaction. This is a feature, not a bug, because it allows you to perform several queries on the statistics and correlate the results without worrying that the numbers are changing underneath you. But if you want to see new results with each query, be sure to do the queries outside any transaction block. Alternatively, you can invoke `pg_stat_clear_snapshot()`, which will discard the current transaction's statistics snapshot (if any). The next use of statistical information will cause a new snapshot to be fetched.

A transaction can also see its own statistics (as yet untransmitted to the collector) in the views `pg_stat_xact_all_tables`, `pg_stat_xact_sys_tables`, `pg_stat_xact_user_tables`, and `pg_stat_xact_user_functions`. These numbers do not act as stated above; instead they update continuously throughout the transaction.

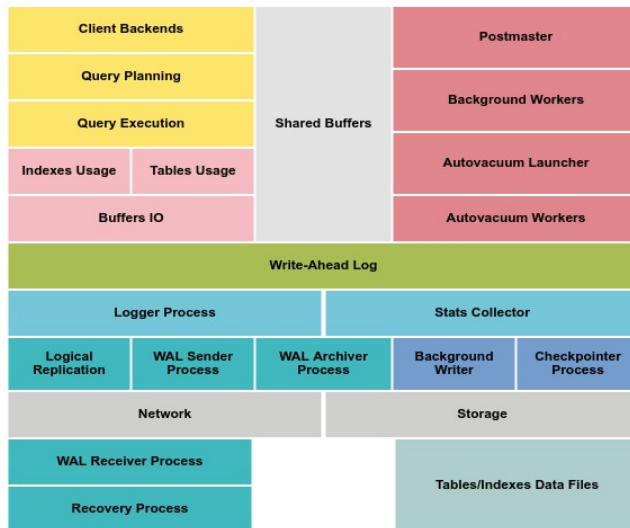
Some of the information in the dynamic statistics views shown in [Table 27.1](#) is security restricted. Ordinary users can only see all the information about their own sessions (sessions belonging to a role that they are a member of). In rows about other sessions, many columns will be null. Note, however, that the existence of a session and its general properties such as its sessions user and database are visible to all users. Superusers and members of the built-in role `pg_read_all_stats` (see also [Section 21.5](#)) can see all the information about all sessions.

Table 27.1. Dynamic Statistics Views

View Name	Description
<code>pg_stat_activity</code>	One row per server process, showing information related to the current activity of that process, such as state and current query. See <a href="#">pg_stat_activity</a> for details.
<code>pg_stat_replication</code>	One row per WAL sender process, showing statistics about replication to that sender's connected standby server. See <a href="#">pg_stat_replication</a> for details.
<code>pg_stat_wal_receiver</code>	Only one row, showing statistics about the WAL receiver from that receiver's connected server. See <a href="#">pg_stat_wal_receiver</a> for details.
<code>pg_stat_subscription</code>	At least one row per subscription, showing information about the subscription workers. See <a href="#">pg_stat_subscription</a> for details.
<code>pg_stat_ssl</code>	One row per connection (regular and replication), showing information about SSL used on this connection. See <a href="#">pg_stat_ssl</a> for details.
<code>pg_stat_ossani</code>	One row per connection (regular and replication), showing information about GSSAPI authentication and encryption used on this connection. See <a href="#">pg_stat_ossani</a> for details.



pg\_stat\_ssl  
 pg\_stat\_activity  
 EXPLAIN  
 pg\_stat\_statements  
 pg\_stat\_user\_functions  
 pg\_prepared\_xacts  
 pg\_locks  
 pg\_stat\_all\_indexes  
 pg\_stat\_all\_tables  
 pg\_statio\_all\_indexes  
 pg\_statio\_all\_tables  
 pg\_statio\_all\_sequences  
 pg\_is\_xlog\_replay\_paused()  
 pg\_current\_xlog\_location()  
 pg\_xlog\_location\_diff()  
 pg\_replication\_slots  
 pg\_stat\_replication  
 pg\_stat\_archiver  
 pg\_stat\_database\_conflicts



pg\_buffercache  
 pg\_stat\_activity  
 pg\_stat\_database  
 pg\_stat\_all\_tables  
 pg\_xlogfile\_name()  
 pg\_current\_xlog\_insert\_location()  
 pg\_xlogfile\_name\_offset()  
 pg\_last\_xlog\_receive\_location()  
 pg\_last\_xlog\_replay\_location()  
 pg\_last\_xact\_replay\_timestamp()  
 pg\_stat\_bgwriter  
 pgstattuple  
 pg\_tablespace\_size()  
 pg\_database\_size()  
 pg\_total\_relation\_size()  
 pg\_relation\_size()  
 pg\_table\_size()  
 pg\_indexes\_size()  
 pg\_is\_dir()  
 pg\_relation\_filenode()  
 pg\_relation\_filepath()  
 pg\_filenode\_relation()

<https://pgstats.dev>







1 / 7

Две минуты  
Турецкий! (с)



*Snatch, 2001*

pg\_stat\_activity — недостаточно

9.6 pg\_stat\_progress\_vacuum

12 pg\_stat\_progress\_create\_index

12 pg\_stat\_progress\_cluster



# pg\_stat\_progress\_analyze, 13

View "pg\_stat\_progress\_analyze"

```
-----  
pid  
datid  
datname  
relid  
phase  
sample_blks_total  
sample_blks_scanned  
ext_stats_total  
ext_stats_computed  
child_tables_total  
child_tables_done  
current_child_table_relid
```



+ этап операции

View "pg\_stat\_progress\_analyze"

```
-----  
pid  
datid  
datname  
relid  
phase  
sample_blks_total  
sample_blks_scanned  
ext_stats_total  
ext_stats_computed  
child_tables_total  
child_tables_done  
current_child_table_relid
```



+ этап операции

+ размер сэмпла и сколько обработано

View "pg\_stat\_progress\_analyze"

```
-----  
pid  
datid  
datname  
relid  
phase  
sample_blks_total  
sample_blks_scanned  
ext_stats_total  
ext_stats_computed  
child_tables_total  
child_tables_done  
current_child_table_relid
```



- + этап операции
- + размер сэмпла и сколько обработано
- + поддержка партиционированных таблиц

View "pg\_stat\_progress\_analyze"

```
-----  
pid  
datid  
datname  
relid  
phase  
sample_blks_total  
sample_blks_scanned  
ext_stats_total  
ext_stats_computed  
child_tables_total  
child_tables_done  
current_child_table_relid
```



- + этап операции
- + размер сэмпла и сколько обработано
- + поддержка партиционированных таблиц
- + соединение с pg\_stat\_activity, pg\_locks

View "pg\_stat\_progress\_analyze"

-----  
pid  
datid  
datname  
relid  
phase  
sample\_blks\_total  
sample\_blks\_scanned  
ext\_stats\_total  
ext\_stats\_computed  
child\_tables\_total  
child\_tables\_done  
current\_child\_table\_relid





# pg\_stat\_progress\_basebackup, 13

pg\_stat\_progress\_basebackup

Columns:

Source:

View pg\_stat\_progress\_basebackup

```
-----  
pid  
phase  
backup_total  
backup_streamed  
tablespaces_total  
tablespaces_streamed
```



+ этап операции

View pg\_stat\_progress\_basebackup

-----  
pid

phase

backup\_total

backup\_streamed

tablespaces\_total

tablespaces\_streamed



+ этап операции

+ полный размер и сколько уже отправлено

View pg\_stat\_progress\_basebackup

```
-----  
pid  
phase  
backup_total  
backup_streamed  
tablespaces_total  
tablespaces_streamed
```



- + этап операции
- + полный размер и сколько уже отправлено
- + соединение с pg\_stat\_activity, pg\_locks?

View pg\_stat\_progress\_basebackup

```
-----  
pid  
phase  
backup_total  
backup_streamed  
tablespaces_total  
tablespaces_streamed
```



# pg\_stat\_progress\_copy, 14

View pg\_stat\_progress\_copy

-----  
pid  
datid  
datname  
relid  
command  
type  
bytes\_processed  
bytes\_total  
tuples\_processed  
tuples\_excluded



+ детали COPY команды

View pg\_stat\_progress\_copy

-----  
pid  
datid  
datname  
relid  
command  
type  
bytes\_processed  
bytes\_total  
tuples\_processed  
tuples\_excluded



+ детали COPY команды

+ сколько обработано байт и строк

View pg\_stat\_progress\_copy

-----  
pid  
datid  
datname  
relid  
command  
type  
bytes\_processed  
bytes\_total  
tuples\_processed  
tuples\_excluded



- + детали COPY команды
- + сколько обработано байт и строк
- + соединение с pg\_stat\_activity, pg\_locks

View pg\_stat\_progress\_copy

-----  
pid  
datid  
datname  
relid  
command  
type  
bytes\_processed  
bytes\_total  
tuples\_processed  
tuples\_excluded





CHECKPOINT

CREATE ..., ALTER ..., DROP ...

REFRESH ...

SELECT ..., INSERT ..., UPDATE ..., DELETE ...



# 2<sub>/7</sub>

**Тик так,  
мистер Уик,  
тик так**



pg\_stat\_activity

- backend\_start
- xact\_start
- query\_start
- state\_change



pg\_stat\_activity

- backend\_start
- xact\_start
- query\_start
- state\_change → state



## pg\_stat\_activity

- backend\_start
- xact\_start → idle
- query\_start → active
- state\_change → state → idle in transaction  
→ idle in transaction (aborted)  
→ fastpath function call



Сколько времени подключения проводят в состояниях?



View pg\_stat\_database

```
-----  
...  
session_time  
active_time  
idle_in_transaction_time  
sessions  
sessions_abandoned  
sessions_fatal  
sessions_killed  
...
```



Время проведенное в сессиях

View pg\_stat\_database

-----  
...  
session\_time  
active\_time  
idle\_in\_transaction\_time  
sessions  
sessions\_abandoned  
sessions\_fatal  
sessions\_killed  
...





Время проведенное в сессиях

Статусы завершения сессий

View pg\_stat\_database

```
-----  
...  
session_time  
active_time  
idle_in_transaction_time  
sessions  
sessions_abandoned  
sessions_fatal  
sessions_killed  
...
```



datname	backends_total	session_total	active_total	idle_xact_total	session,ms	active,ms	idle_xact,ms	sessions	abandoned	fatal	killed
lesovsky	1	166:09:30	00:20:37	00:00:00	1012.78	11.84	0.00	0	0	0	0
pgbench	0	122:11:18	10:28:28	00:16:57	0.00	0.00	0.00	0	0	0	0
postgres	0	05:25:01	00:42:53	00:00:00	0.00	0.00	0	0	0	0	0
flight_v2	0	01:01:21	00:06:23	00:00:00	0.00	0.00	0.00	0	0	0	0
flight	1	00:32:47	00:06:14	00:00:00	1692.82	322.24	0.00	698	0	0	0
env_signatures	0	00:22:31	00:20:05	00:00:06	0.00	0.00	0.00	0	0	0	0
pgscv_fixtures	0	00:22:29	00:17:00	00:00:00	0.00	0.00	0	0	0	0	0
template0	0	00:00:00	00:00:00	00:00:00	0	0	0	0	0	0	0
	0	00:00:00	00:00:00	00:00:00	0	0	0	0	0	0	0
template1	0	00:00:00	00:00:00	00:00:00	0.00	0.00	0	0	0	0	0



# 3<sub>/7</sub>

**Nobody panics  
when things go  
according to plan.  
Even if the plan is  
horrifying!**



*The Dark Knight. 2008*

EXPLAIN и auto\_explain — выборочное планирование

- Ручной запуск (EXPLAIN)
- Поиск и анализ журналов



## Время планирования в pg\_stat\_statements (13)

- time переименовали в exec\_time
- добавили plan\_time
- pg\_stat\_statements.track\_planning -- default: off



## Отметка уровня выполнения (14)

- `pg_stat_statements.track`: all vs top
- `track = all` — двойной подсчет в агрегатах
- `pg_stat_statements.toplevel`



Track rows for utility commands (14)

- REFRESH, CREATE TABLE AS, SELECT INTO, FETCH

pg\_stat\_statements\_info (14) — dealloc, stats\_reset

Track WAL usage stats (13)



# 4<sub>7</sub>

**У меня, все  
ходы  
записаны! (с)**



12 стульев, 1976



Объем записи от запросов? `pg_stat_statements`

Объем записи от фоновых служб? `pg_stat_bgwriter`

Объем записи WAL — ?



Объем записи от запросов? `pg_stat_statements`

Объем записи от фоновых служб? `pg_stat_bgwriter`

Объем записи WAL — `SELECT pg_current_wal_lsn() - '0/0'::pg_lsn;`



View pg\_stat\_wal

```
-----  
wal_records  
wal_fpi  
wal_bytes  
wal_buffers_full  
wal_write  
wal_sync  
wal_write_time  
wal_sync_time  
stats_reset
```



## Количественные характеристики

View pg\_stat\_wal

```
-----  
wal_records  
wal_fpi  
wal_bytes  
wal_buffers_full  
wal_write  
wal_sync  
wal_write_time  
wal_sync_time  
stats_reset
```



Количественные характеристики

Нехватка wal\_buffers

```
View pg_stat_wal
-----
wal_records
wal_fpi
wal_bytes
wal_buffers_full
wal_write
wal_sync
wal_write_time
wal_sync_time
stats_reset
```



Количественные характеристики

Нехватка wal\_buffers

Затраченное время

```
View pg_stat_wal
```

```
-----  
wal_records  
wal_fpi  
wal_bytes  
wal_buffers_full  
wal_write  
wal_sync  
wal_write_time  
wal_sync_time  
stats_reset
```



Трекинг WAL статистики, 13:

- EXPLAIN, auto\_explain
- autovacuum
- pg\_stat\_statements



**5<sub>/7</sub>**

**One does not  
simply (c)**



*The Lord of the Rings: The Fellowship of the Ring, 2001*



Комплексный учет памяти: RES, VIRT, SHM, HUGE\_PAGES и т.п.

top, htop могут, и часто вводят в заблуждение

Можно использовать данные из /proc/\$PID/

<https://blog.anarazel.de/2020/10/07/measuring-the-memory-overhead-of-a-postgres-connection/>



# pg\_backend\_memory\_contexts, 14

View pg\_backend\_memory\_contexts

-----  
name  
ident  
parent  
level  
total\_bytes  
total\_nblocks  
free\_bytes  
free\_chunks  
used\_bytes



Детали об участке памяти

View pg\_backend\_memory\_contexts

-----  
name  
ident  
parent  
level  
total\_bytes  
total\_nblocks  
free\_bytes  
free\_chunks  
used\_bytes



Детали об участке памяти

Статистика total, free, used

View pg\_backend\_memory\_contexts

-----  
name  
ident  
parent  
level  
total\_bytes  
total\_nblocks  
free\_bytes  
free\_chunks  
used\_bytes



Детали об участке памяти

Статистика total, free, used

`pg_log_backend_memory_contexts(integer)`



# 6<sub>/7</sub>

- Dance with me.
- You don't dance.
- It was just my cover.
- Was sloth your cover, too?



pg\_stat\_activity

pg\_stat\_statements



pg\_stat\_activity.queryid

pg\_stat\_statements.queryid

compute\_query\_id, 14





```
select a.*, s.* from pg_stat_activity a inner join pg_stat_statements s on (a.datid = s.dbid AND a.usesysid = s.userid AND a.query_id = s.queryid)
where a.pid = 1001291;
```

```
-[ RECORD 1 ]-----+-----
datid          | 16413
datname        | pgbench
pid            | 1001291
leader_pid     |
usesysid       | 10
username       | postgres
application_name | pgbench
client_addr    |
client_hostname |
client_port    | -1
backend_start  | 2021-05-22 10:15:57.299468+05
xact_start     | 2021-05-22 10:16:25.566151+05
query_start    | 2021-05-22 10:16:25.566623+05
state_change   | 2021-05-22 10:16:25.566763+05
wait_event_type |
wait_event     |
state          | idle in transaction
backend_xid     | 237577
backend_xmin    |
query_id       | 2517686606037258902
query          | SELECT abalance FROM pgbench_accounts WHERE aid = 1715456;
backend_type    | client backend
userid         | 10
dbid           | 16413
toplevel       | t
queryid        | 2517686606037258902
query          | SELECT abalance FROM pgbench_accounts WHERE aid = $1
plans          | 0
total_plan_time | 0
min_plan_time   | 0
max_plan_time   | 0
mean_plan_time  | 0
stddev_plan_time | 0
calls          | 209439
total_exec_time | 4251.98569499987
min_exec_time   | 0.005414
max_exec_time   | 0.435581
mean_exec_time  | 0.020301785698938563
stddev_exec_time | 0.005889254053319066
rows           | 209439
shared_blks_hit | 884097
```

pg\_stat\_activity.leader\_pid (13)



```
lesovsky@matanuii:~ $ pgcenter profile -P 1050881 -W -U postgres pgbench
```

```
LOG: Profiling process 1050881 with 100ms sampling
```

```
-----
% time      seconds wait_event      query: select count(*) from pgbench_accounts ;
-----
90.91      61.090423 IO.DataFileRead
 7.74       5.203075 Running
 1.34       0.902708 IPC.BtreePage
-----
100.00      67.196206 including workers
              67.196206
```

```
^Cgot interrupt
```

```
lesovsky@matanuii:~ $ pgcenter profile -P 1050881 -U postgres pgbench
```

```
LOG: Profiling process 1050881 with 100ms sampling
```

```
-----
% time      seconds wait_event      query: select count(*) from pgbench_accounts ;
-----
90.00      176.637457 IO.DataFileRead
 9.73       19.102891 Running
 0.20       0.401478 IPC.BtreePage
 0.06       0.121435 LWLock.BufferIO
-----
100.00      196.263261 including workers
              65.421091
```

```
^Cgot interrupt
```

```
lesovsky@matanuii:~ $ █
```



7<sub>17</sub>



`pg_shmem_allocations` — распределение памяти в shared buffers (13)



pg\_shmem\_allocations — распределение памяти в shared buffers (13)

pg\_stat\_slru — использование SLRU кэшей (13)



pg\_shmem\_allocations — распределение памяти в shared buffers (13)

pg\_stat\_slru — использование SLRU кэшей (13)

pg\_stat\_replication\_slots — статистика слотов репликации (14)



pg\_shmem\_allocations — распределение памяти в shared buffers (13)

pg\_stat\_slru — использование SLRU кэшей (13)

pg\_stat\_replication\_slots — статистика слотов репликации (14)

pg\_locks.waitstart (14)





pg\_shmem\_allocations — распределение памяти в shared buffers (13)

pg\_stat\_slru — использование SLRU кэшей (13)

pg\_stat\_replication\_slots — статистика слотов репликации (14)

pg\_locks.waitstart (14)

Время выполнения IO при логировании **auto-vacuum/analyze** (14)



THE END

pg\_stat\_activity, pg\_stat\_progress\_\*

pg\_stat\_statements







## pg\_stat\_activity

## pg\_stat\_activity

Related items: [Client Backends](#), [Query Execution](#), [Background Workers](#), [Autovacuum Launcher](#), [Autovacuum Workers](#)

The `pg_stat_activity` view will have one row per server process, showing information related to the current activity of that process.

## TIPS'N'TRICKS

- Use `now()` - `xact_start` for get human-readable age of transactions. Same approach can be used with `'backend_start'`, `'query_start'` and `'state_change'`.
- Use `'coalesce(username, backend_type)'`, because system processes like `autovacuum` workers have `NULL` in `'username'`.

## RESOURCES

- [Official documentation](#)
- [Summary activity](#). Brief information about connected clients.
- [Long activity](#). Current activity longer than 100ms.

## CHANGES

- 14: Add 'query\_id' to report query ID of recent query executed by backend.
- 13: Add 'leader\_pid' to report a parallel worker's leader process.
- 10: Add 'backend\_type' column which shows auxiliary processes, background workers, and wal sender processes.
- 10: Allow to show the SQL query being executed by parallel workers.
- 9.6: Replace the 'waiting' column with 'wait\_event\_type' and 'wait\_event'.

### COLUMNS DESCRIPTION

name	type	description
datid	oid	OID of the database this backend is connected to.
datname	name	Name of the database this backend is connected to.
pid	integer	Process ID of this backend.
leader_pid	integer	Process ID of the parallel group leader, if this process is a parallel query worker. NULL if this process is a parallel group leader or does not participate in parallel query. <b>Added in Postgres 13.</b>
usesysid	oid	OID of the user logged into this backend.
username	name	Name of the user logged into this backend.
application_name	text	Name of the application that is connected to this backend.
client_addr	inet	IP address of the client connected to this backend. If this field is NULL, it indicates either that the client is connected via a Unix socket on the server machine or that this is an internal process such as autovacuum.
client_hostname	text	Host name of the connected client, as reported by a reverse DNS lookup of client_addr. This field will only be non-NULL for IP connections, and only when 'log_hostname' is enabled.
client_port	integer	TCP port number that the client is using for communication with this backend, or -1 if a Unix socket is used. If this field is NULL, it indicates that this is an internal server process.
backend_start	timestamp with time zone	Time when this process was started. For client backends, this is the time the client connected to the server.
xact_start	timestamp with time zone	Time when this process' current transaction was started, or NULL if no transaction is active. If the current query is the first of its transaction, this column is equal to the 'query_start' column.
query_start	timestamp with time zone	Time when the currently active query was started, or if 'state' is not active, when the last query was started.
state_change	timestamp with time zone	Time when the state was last changed.
waiting	boolean	True if this backend is currently waiting on a lock. <b>Removed in Postgres 9.6.</b>
wait_event_type	text	The type of event for which the backend is waiting, if any; otherwise NULL. See <a href="#">Wait Event Types</a> table. <b>Added in Postgres 9.6.</b>
wait_event	text	Wait event name if backend is currently waiting, otherwise NULL. See the group of <a href="#">Wait Events</a> tables. <b>Added in Postgres 9.6.</b>
		Current overall state of this backend. Possible values are: * active: The backend is executing a query. * idle: The backend is waiting for a new client command. * idle in transaction: The backend is in a transaction, but is not



**Спасибо за внимание**