

# PostgreSQL Buffers

*Vik Fearing*

PGConf.EU Warsaw  
October 25, 2017

# 2ndQuadrant<sup>®</sup> PostgreSQL


**Vik Fearing**  
*Change “and and” to “and”.*

irc: xocolatl


Twitter: @pg\_xocolatl

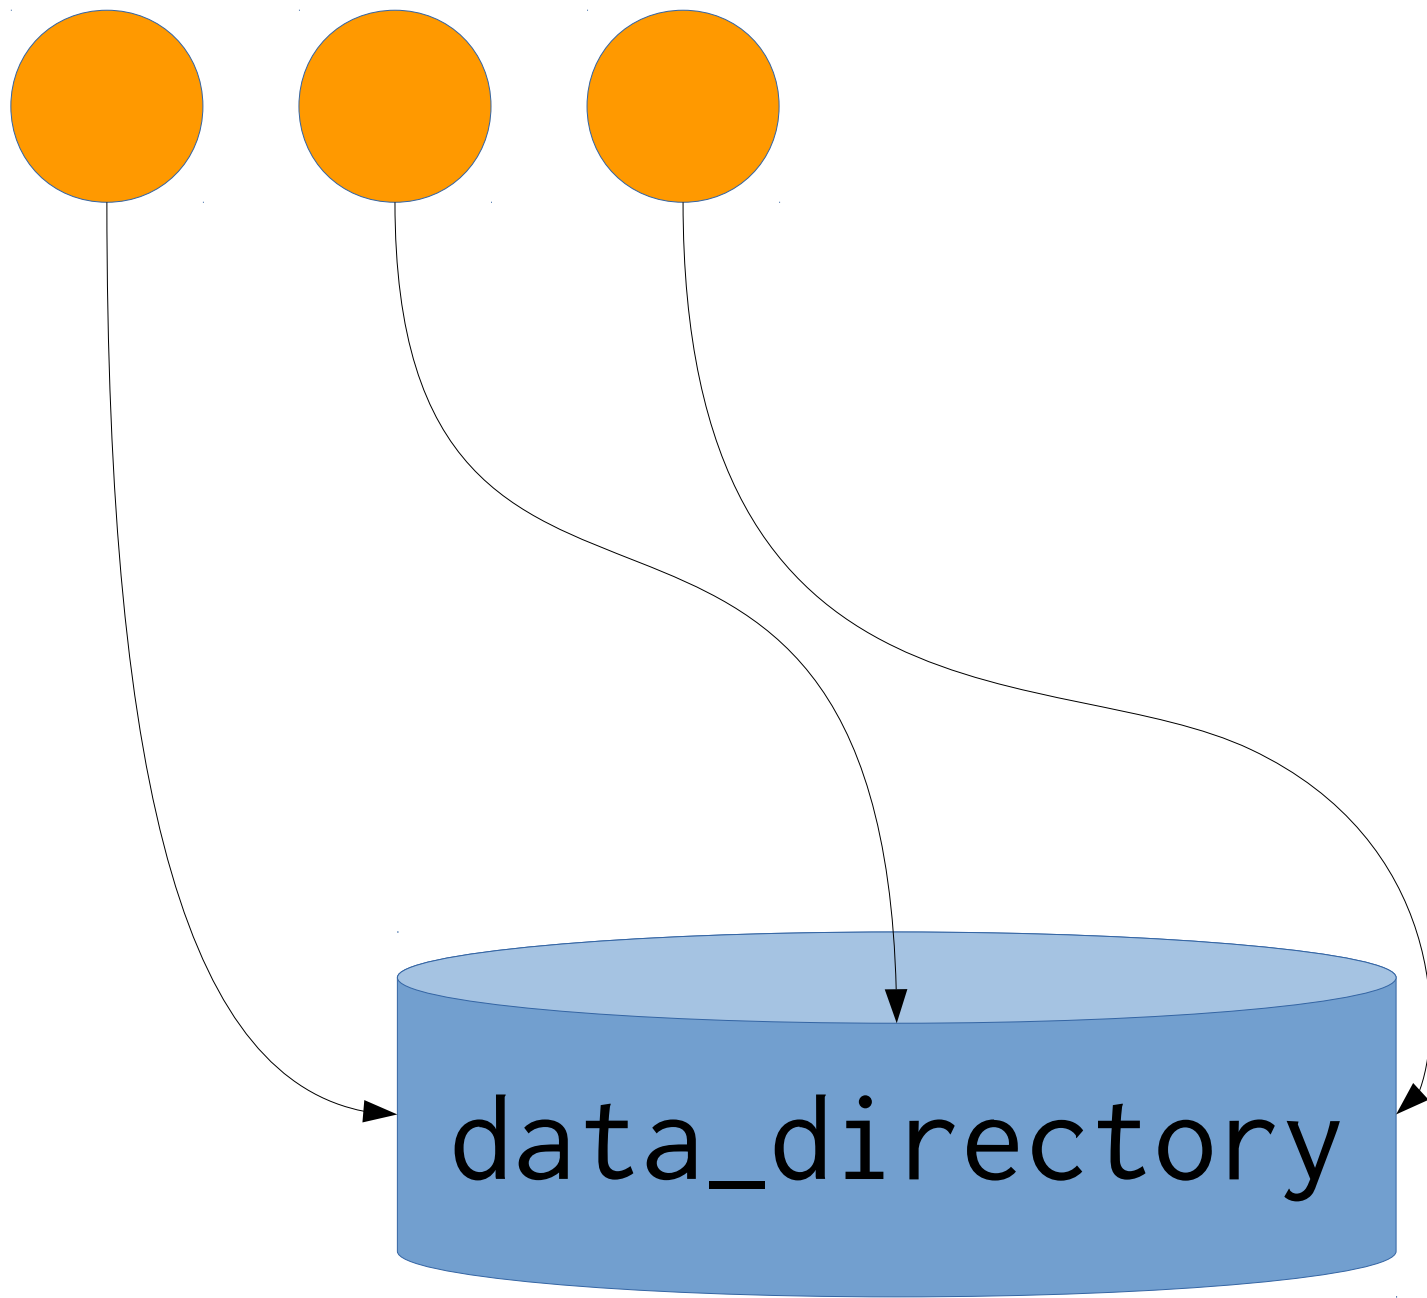
# PostgreSQL Buffers


# Backend



# Data Directory







**Page**

=

**Block**

=

**Buffer**

**8kB**


BLOCKSZ



# Shared Buffers



shared\_buffers



# Usage Count


1 .. 5

**ClockSweep**





Whups!

# Write-Ahead Log (WAL)



# Write-Ahead Log (WAL)








# WAL Buffers

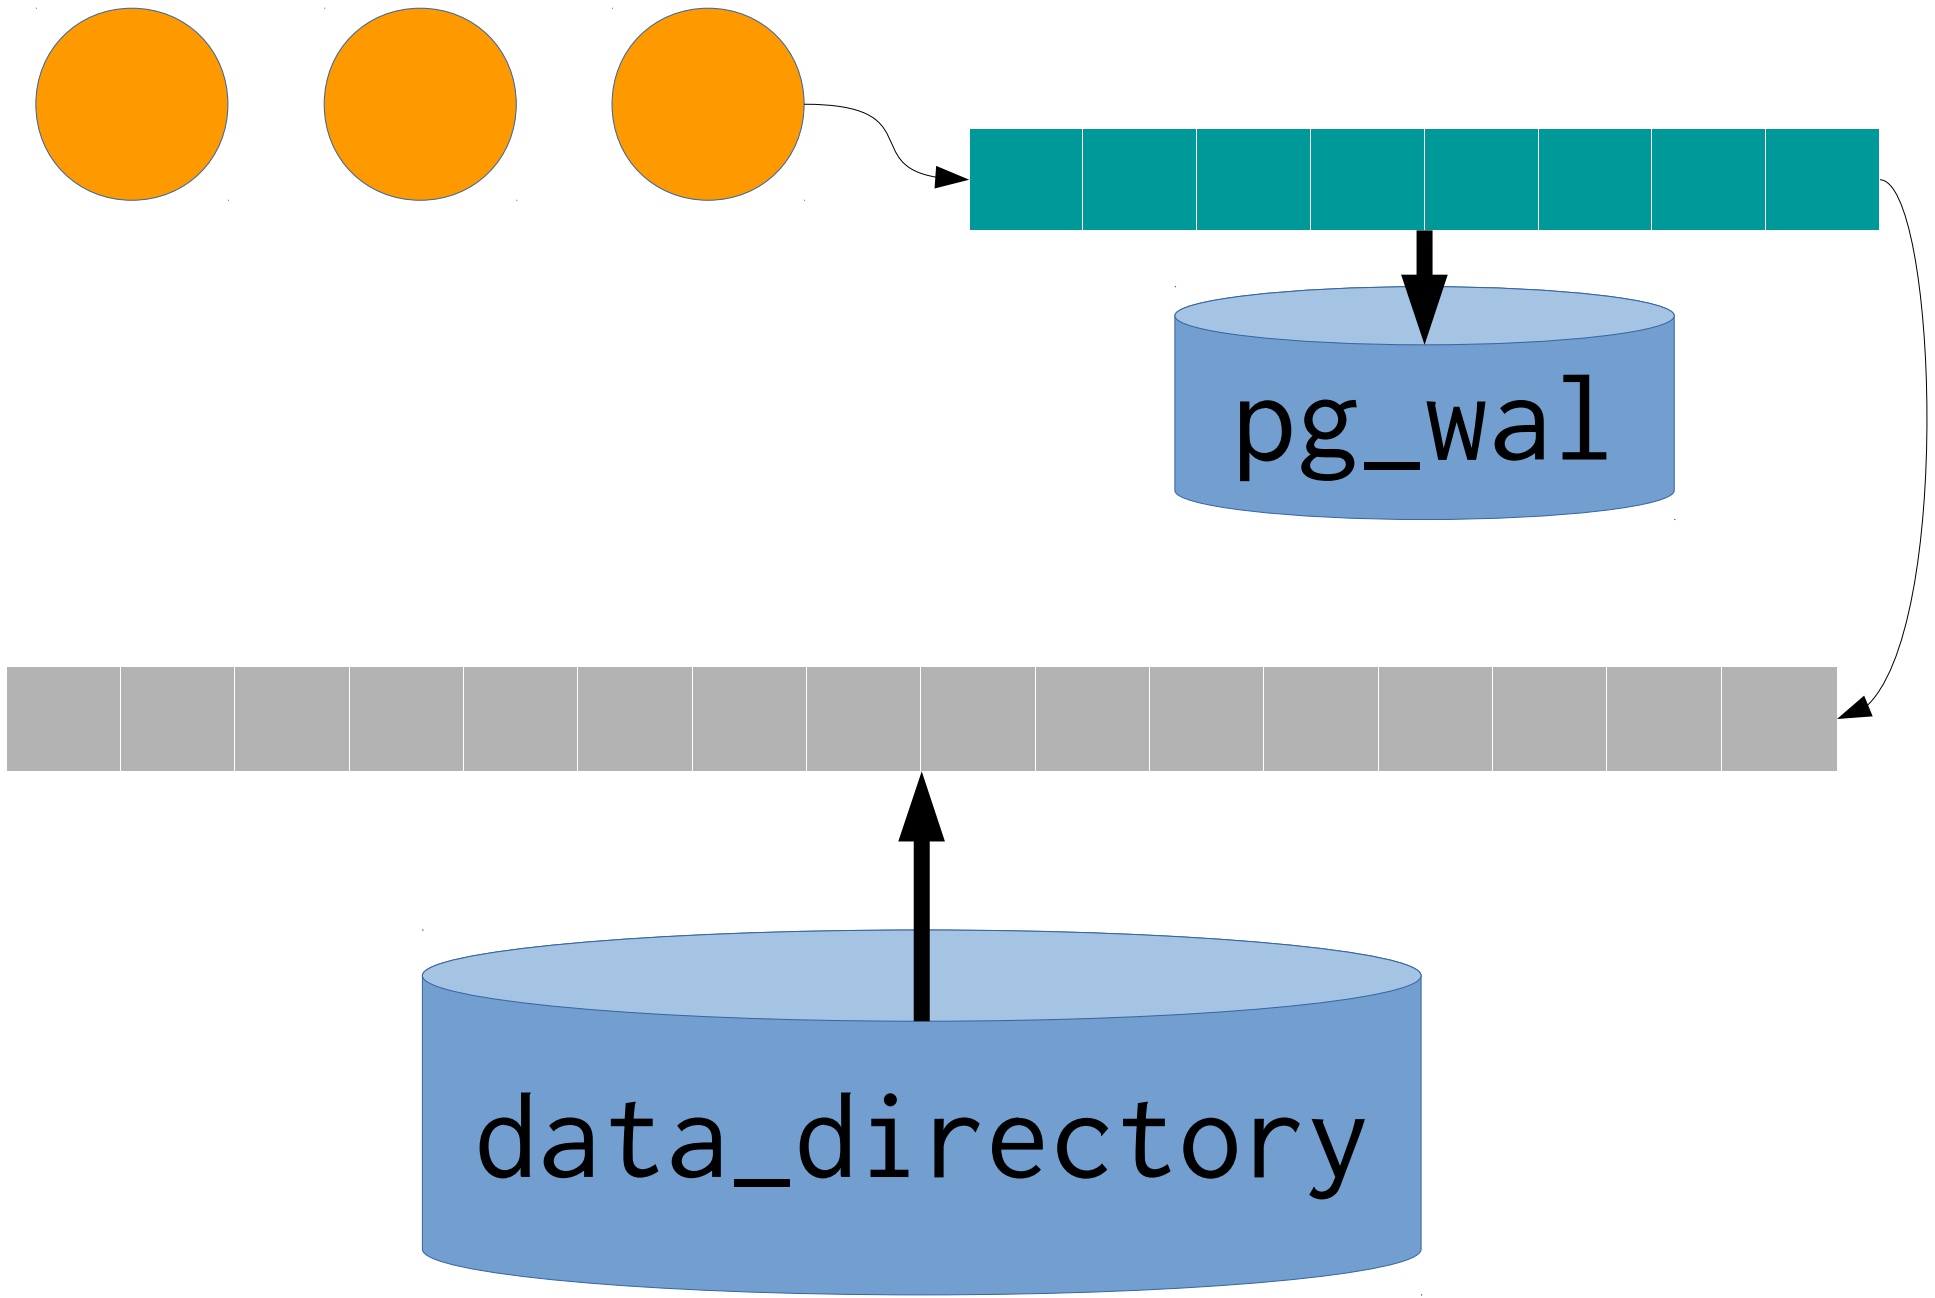
wal\_buffers





# **WAL Writer**

wal\_writer\_delay



# Background Writer


bgwriter\_lru\_multiplier  
bgwriter\_lru\_maxpages  
bgwriter\_delay

# Checkpoints

checkpoint\_timeout

checkpoint\_completion\_target

min\_wal\_size • max\_wal\_size









# Temporary Tables

temp\_buffers

temp\_tablespaces

~~temp\_file\_limit~~





View "public.pg\_buffercache"

Column	Type	Modifiers
bufferid	integer	
relfilenode	oid	
reltablespace	oid	
reldatabase	oid	
relforknumber	smallint	
relblocknumber	bigint	
isdirty	boolean	
usagecount	smallint	
pinning_backends	integer	

## View "pg\_catalog.pg\_stat\_bgwriter"

Column	Type	Modifiers
checkpoints_timed	bigint	
checkpoints_req	bigint	
checkpoint_write_time	double precision	
checkpoint_sync_time	double precision	
buffers_checkpoint	bigint	
buffers_clean	bigint	
maxwritten_clean	bigint	
buffers_backend	bigint	
buffers_backend_fsync	bigint	
buffers_alloc	bigint	
stats_reset	timestamp with time zone	

```
SET track_io_timing = on;
```

```
EXPLAIN (ANALYZE, BUFFERS)
```

-> Nested Loop Left Join  
    (cost=365718.92..9633816.57 rows=1270632 width=84)  
    (actual time=31305.214..75833773.678 rows=2175641 loops=1)  
Join Filter: (u.x = v.y)  
Rows Removed by Join Filter: 18363801709  
Buffers: shared hit=6003348712 read=2378746 dirtied=7984  
I/O Timings: read=961454.651



## View "pg\_catalog.pg\_stat\_database"

Column	Type	Modifiers
datid	oid	
datname	name	
numbackends	integer	
xact_commit	bigint	
xact_rollback	bigint	
blks_read	bigint	
blks_hit	bigint	
tup_returned	bigint	
tup_fetched	bigint	
tup_inserted	bigint	
tup_updated	bigint	
tup_deleted	bigint	
conflicts	bigint	
temp_files	bigint	
temp_bytes	bigint	
deadlocks	bigint	
blk_read_time	double precision	
blk_write_time	double precision	
stats_reset	timestamp with time zone	

View "pg\_catalog.pg\_statio\_all\_tables"

Column	Type	Modifiers
relid	oid	
schemaname	name	
relname	name	
heap_blks_read	bigint	
heap_blks_hit	bigint	
idx_blks_read	bigint	
idx_blks_hit	bigint	
toast_blks_read	bigint	
toast_blks_hit	bigint	
tidx_blks_read	bigint	
tidx_blks_hit	bigint	

View "pg\_catalog.pg\_statio\_all\_indexes"

Column	Type	Modifiers
relid	oid	
indexrelid	oid	
schemaname	name	
relname	name	
indexrelname	name	
idx_blks_read	bigint	
idx_blks_hit	bigint	

## View "public.pg\_stat\_statements"

Column	Type	Modifiers
userid	oid	
dbid	oid	
queryid	bigint	
query	text	
calls	bigint	
total_time	double precision	
min_time	double precision	
max_time	double precision	
mean_time	double precision	
stddev_time	double precision	
rows	bigint	
shared_blks_hit	bigint	
shared_blks_read	bigint	
shared_blks_dirtied	bigint	
shared_blks_written	bigint	
local_blks_hit	bigint	
local_blks_read	bigint	
local_blks_dirtied	bigint	
local_blks_written	bigint	
temp_blks_read	bigint	
temp_blks_written	bigint	
blk_read_time	double precision	
blk_write_time	double precision	

# 2ndQuadrant<sup>®</sup> PostgreSQL

**Vik Fearing**  
*Change “and and” to “and”.*

irc: xocolatl

Twitter: @pg\_xocolatl