

BUFFERS

by default

Michael Christofides

Hi, I'm Michael

I run pgMustard, a tool for explain (~5 years)

And co-host postgres.fm, a podcast (~1.5 years)

pgMustard

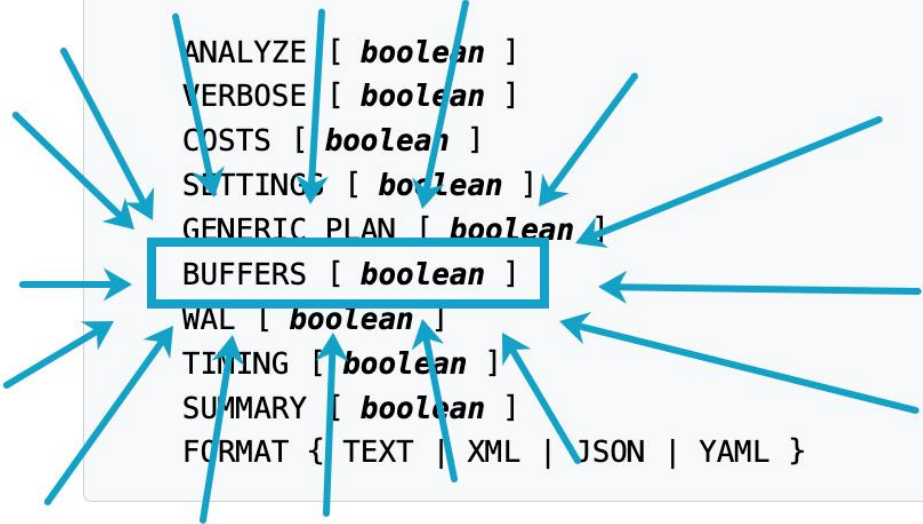
postgres.fm

# What are BUFFERS?

```
EXPLAIN [ ( option [, ...] ) ] statement  
EXPLAIN [ ANALYZE ] [ VERBOSE ] statement
```

where *option* can be one of:

```
ANALYZE [ boolean ]  
VERBOSE [ boolean ]  
COSTS [ boolean ]  
SETTINGS [ boolean ]  
GENERIC PLAN [ boolean ]  
BUFFERS [ boolean ]  
WAL [ boolean ]  
TIMING [ boolean ]  
SUMMARY [ boolean ]  
FORMAT { TEXT | XML | JSON | YAML }
```

A diagram consisting of several blue arrows pointing towards the 'BUFFERS [ boolean ]' option in the list of options. The 'BUFFERS' option is highlighted with a blue rectangular border. The arrows originate from various directions around the text and converge on the 'BUFFERS' option.



`explain analyze`



`explain (analyze, buffers)`

# What are BUFFERS?

Shows details about **the data read/written**

- \* Number of data pages

- \* 8 KB per page

# What are BUFFERS?

Ten statistics, made up of two halves:

1. **Shared** / Local / Temp

2. **Hit** / **Read** / Written / Dirtied

When should you use them?

- \* Performance investigations
- \* Education – blog posts, presentations, ...
- \* When (small) overhead not an issue

**Or, almost always!**

Exhibit A



# Exhibit A

Index Scan using idx on people (cost=0.43..8.45 rows=1  
width=17) (actual time=0.917..0.920 rows=1 loops=1)

Index Cond: (id = 99999)

Buffers: <redacted>

Planning Time: 0.252 ms

**Execution Time: 0.955 ms**

# Exhibit A

Index Scan using idx on people (cost=0.43..8.45 rows=1  
width=17) (actual time=0.079..0.080 rows=1 loops=1)

Index Cond: (id = 99999)

Buffers: <redacted>

Planning Time: 0.110 ms

**Execution Time: 0.111 ms**

# Exhibit A

Index Scan using idx on people (cost=0.43..8.45 rows=1  
width=17) (actual time=0.917..0.920 rows=1 loops=1)

Index Cond: (id = 99999)

**Buffers: shared hit=1 read=3**

Planning Time: 0.252 ms

**Execution Time: 0.955 ms**

# Exhibit A

Index Scan using idx on people (cost=0.43..8.45 rows=1  
width=17) (actual time=0.079..0.080 rows=1 loops=1)

Index Cond: (id = 99999)

**Buffers: shared hit=4**

Planning Time: 0.110 ms

**Execution Time: 0.111 ms**

Exhibit B

# Exhibit B

Aggregate (cost=198437.31..198437.32 rows=1 width=8) (actual time=748.648..748.648 rows=1 loops=1)

-> Seq Scan on people (cost=0.00..173436.25 rows=10000425 width=0) (actual time=0.056..500.741 rows=10000000 loops=1)

Planning Time: 0.166 ms

Execution Time: 748.739 ms

# Exhibit B

Aggregate (cost=198437.31..198437.32 rows=1 width=8) (actual time=712.345..712.346 rows=1 loops=1)

**Buffers: shared hit=16129 read=57303**

-> Seq Scan on people (cost=0.00..173436.25 rows=10000425 width=0) (actual time=0.018..457.055 rows=10000000 loops=1)

**Buffers: shared hit=16129 read=57303**

Planning Time: 0.085 ms

Execution Time: 712.378 ms

$(16129 + 57303) * 8KB \approx 600MB$

# Re-cap

- \* BUFFERS show data **pages read/written**
- \* Helps to show **why** things are slow
- \* Convert to **bytes** for bonus points
- \* explain (analyze, buffers) > explain analyze



# Learn more

- \* [EXPLAIN parameters](#) (docs)
- \* [EXPLAIN glossary – buffers](#) (my site)
- \* [Buffers by default](#) (podcast episode)
- \* [Latest patch to add buffers by default](#) 🙌

Slides: [pgmustard.com/paris24](http://pgmustard.com/paris24)

Fin