

A hand is shown typing on a laptop keyboard, with a blue overlay and several white rounded squares scattered across the scene. The text is overlaid on the left side of the image.

# **Acelerando a Manada**

## **Parallel Queries no PostgreSQL**

# Paralelismo Intra-Consulta em PostgreSQL pré 9.6



- Apresenta maiores ganhos em grandes volumes (Pode não encaixar em índices muito seletivos)
- Execução sequencial pode ser mais rápida
- Limitações de implementação
  - Consultas que contém operações de modificações de dados
  - Consultas que possam ser suspensas (Cursors)
  - Subconsulta de uma consulta já paralelizada
  - Nível de Isolamento de Transação = Serializable

- `max_worker_processes` (8)
- `max_parallel_workers_per_gather` (2)
- `max_parallel_workers` (8) \*
- `dynamic_shared_memory_type` (posix)

- Paralelismo inicia com Busca Sequencial
- Parallel Join
  - Nested Loop
  - Hash Join
- Não disponível para índices

```
cliente=> EXPLAIN ANALYZE SELECT COUNT(*) FROM documento;
```

```
Aggregate (cost=127534.05..127534.06 rows=1 width=8) (actual time=564.417..564.418 rows=1 loops=1)  
  -> Seq Scan on documento (cost=0.00..122195.84 rows=2135284 width=0) (actual time=0.037..431.728 rows=2135284 loops=1)  
Planning time: 0.207 ms  
Execution time: 564.551 ms  
(4 registros)
```

```
Aggregate (cost=127534.05..127534.06 rows=1 width=8) (actual time=544.880..544.880 rows=1 loops=1)  
  -> Seq Scan on documento (cost=0.00..122195.84 rows=2135284 width=0) (actual time=0.041..405.056 rows=2135284 loops=1)  
Planning time: 0.208 ms  
Execution time: 544.952 ms  
(4 registros)
```

```
Aggregate (cost=127534.05..127534.06 rows=1 width=8) (actual time=546.038..546.038 rows=1 loops=1)  
  -> Seq Scan on documento (cost=0.00..122195.84 rows=2135284 width=0) (actual time=0.033..406.194 rows=2135284 loops=1)  
Planning time: 0.210 ms  
Execution time: 546.109 ms  
(4 registros)
```

```
cliente=> EXPLAIN ANALYZE SELECT COUNT(*) FROM documento;
```

```
Finalize Aggregate (cost=107181.73..107181.74 rows=1 width=8) (actual time=125.449..125.449 rows=1 loops=1)
-> Gather (cost=107181.21..107181.72 rows=5 width=8) (actual time=125.397..125.443 rows=6 loops=1)
    Workers Planned: 5
    Workers Launched: 5
-> Partial Aggregate (cost=106181.21..106181.22 rows=1 width=8) (actual time=121.325..121.325 rows=1 loops=6)
    -> Parallel Seq Scan on documento (cost=0.00..105113.57 rows=427057 width=0) (actual time=0.033..98.871
rows=355881 loops=6)
    Planning time: 0.216 ms
    Execution time: 129.150 ms
(8 registros)
```

```
Finalize Aggregate (cost=107181.73..107181.74 rows=1 width=8) (actual time=128.315..128.315 rows=1 loops=1)
-> Gather (cost=107181.21..107181.72 rows=5 width=8) (actual time=128.254..128.307 rows=6 loops=1)
    Workers Planned: 5
    Workers Launched: 5
-> Partial Aggregate (cost=106181.21..106181.22 rows=1 width=8) (actual time=124.265..124.265 rows=1 loops=6)
    -> Parallel Seq Scan on documento (cost=0.00..105113.57 rows=427057 width=0) (actual time=0.037..101.476
rows=355881 loops=6)
    Planning time: 0.217 ms
    Execution time: 132.337 ms
(8 registros)
```

```
cliente=> EXPLAIN ANALYZE SELECT COUNT(*) FROM documento WHERE agenciaID = 36;
```

QUERY PLAN

```
-----  
Aggregate (cost=127910.40..127910.40 rows=1 width=8) (actual time=504.106..504.106 rows=1 loops=1)  
-> Seq Scan on documento (cost=0.00..127534.05 rows=150538 width=0) (actual time=0.035..491.605 rows=153491 loops=1)  
    Filter: (agenciaID = 36)  
    Rows Removed by Filter: 1981793  
Planning time: 0.176 ms  
Execution time: 504.168 ms  
(6 registros)
```

```
Aggregate (cost=127910.40..127910.40 rows=1 width=8) (actual time=506.751..506.751 rows=1 loops=1)  
-> Seq Scan on documento (cost=0.00..127534.05 rows=150538 width=0) (actual time=0.046..494.060 rows=153491 loops=1)  
    Filter: (agenciaID = 36)  
    Rows Removed by Filter: 1981793  
Planning time: 0.188 ms  
Execution time: 506.817 ms
```

```
Aggregate (cost=127910.40..127910.40 rows=1 width=8) (actual time=513.973..513.973 rows=1 loops=1)  
-> Seq Scan on documento (cost=0.00..127534.05 rows=150538 width=0) (actual time=0.040..501.330 rows=153491 loops=1)  
    Filter: (agenciaID = 36)  
    Rows Removed by Filter: 1981793  
Planning time: 0.185 ms  
Execution time: 514.036 ms
```



```
cliente=> EXPLAIN ANALYZE SELECT COUNT(*) FROM documento WHERE agenciaID = 36;
```

```
Finalize Aggregate (cost=107257.00..107257.01 rows=1 width=8) (actual time=103.679..103.679 rows=1 loops=1)
-> Gather (cost=107256.48..107256.99 rows=5 width=8) (actual time=103.623..103.673 rows=6 loops=1)
    Workers Planned: 5
    Workers Launched: 5
-> Partial Aggregate (cost=106256.48..106256.49 rows=1 width=8) (actual time=100.722..100.722 rows=1 loops=6)
    -> Parallel Seq Scan on documento (cost=0.00..106181.21 rows=30108 width=0) (actual time=0.052..98.624
rows=25582 loops=6)
        Filter: (agenciaID = 36)
        Rows Removed by Filter: 330299
Planning time: 0.141 ms
Execution time: 107.316 ms
```

```
Finalize Aggregate (cost=107257.00..107257.01 rows=1 width=8) (actual time=120.814..120.814 rows=1 loops=1)
-> Gather (cost=107256.48..107256.99 rows=5 width=8) (actual time=120.757..120.806 rows=6 loops=1)
    Workers Planned: 5
    Workers Launched: 5
-> Partial Aggregate (cost=106256.48..106256.49 rows=1 width=8) (actual time=116.780..116.780 rows=1 loops=6)
    -> Parallel Seq Scan on documento (cost=0.00..106181.21 rows=30108 width=0) (actual time=0.097..114.617
rows=25582 loops=6)
        Filter: (agenciaID = 36)
        Rows Removed by Filter: 330299
Planning time: 0.200 ms
Execution time: 124.804 ms
```

```
Finalize Aggregate (cost=107257.00..107257.01 rows=1 width=8) (actual time=117.424..117.424 rows=1 loops=1)
-> Gather (cost=107256.48..107256.99 rows=5 width=8) (actual time=117.366..117.418 rows=6 loops=1)
    Workers Planned: 5
    Workers Launched: 5
-> Partial Aggregate (cost=106256.48..106256.49 rows=1 width=8) (actual time=113.459..113.459 rows=1 loops=6)
    -> Parallel Seq Scan on documento (cost=0.00..106181.21 rows=30108 width=0) (actual time=0.050..111.285
rows=25582 loops=6)
        Filter: (agenciaID = 36)
        Rows Removed by Filter: 330299
Planning time: 0.196 ms
Execution time: 121.447 ms
```

```
cliente=> EXPLAIN ANALYZE SELECT COUNT(*) FROM documento JOIN agencia USING (agenciaID);
```

```
Finalize Aggregate (cost=113059.66..113059.67 rows=1 width=8) (actual time=198.927..198.927 rows=1 loops=1)
-> Gather (cost=113059.13..113059.64 rows=5 width=8) (actual time=198.865..198.919 rows=6 loops=1)
    Workers Planned: 5
    Workers Launched: 5
-> Partial Aggregate (cost=112059.13..112059.14 rows=1 width=8) (actual time=194.367..194.367 rows=1 loops=6)
    -> Hash Join (cost=7.60..110991.49 rows=427057 width=0) (actual time=0.236..171.281 rows=355763 loops=6)
        Hash Cond: (documento.agenciaID = agencia.agenciaID)
        -> Parallel Seq Scan on documento (cost=0.00..105113.57 rows=427057 width=4) (actual
time=0.037..90.926 rows=355881 loops=6)
        -> Hash (cost=5.60..5.60 rows=160 width=4) (actual time=0.073..0.073 rows=160 loops=6)
            Buckets: 1024 Batches: 1 Memory Usage: 14kB
            -> Seq Scan on agencia (cost=0.00..5.60 rows=160 width=4) (actual time=0.017..0.047 rows=160
loops=6)
    Planning time: 1.564 ms
    Execution time: 202.854 ms
(13 registros)
```

```
Finalize Aggregate (cost=113059.66..113059.67 rows=1 width=8) (actual time=189.303..189.303 rows=1 loops=1)
-> Gather (cost=113059.13..113059.64 rows=5 width=8) (actual time=189.212..189.297 rows=6 loops=1)
    Workers Planned: 5
    Workers Launched: 5
-> Partial Aggregate (cost=112059.13..112059.14 rows=1 width=8) (actual time=185.829..185.829 rows=1 loops=6)
    -> Hash Join (cost=7.60..110991.49 rows=427057 width=0) (actual time=0.200..163.117 rows=355763 loops=6)
        Hash Cond: (documento.agenciaID = agencia.agenciaID)
        -> Parallel Seq Scan on documento (cost=0.00..105113.57 rows=427057 width=4) (actual
time=0.027..82.541 rows=355881 loops=6)
        -> Hash (cost=5.60..5.60 rows=160 width=4) (actual time=0.066..0.066 rows=160 loops=6)
            Buckets: 1024 Batches: 1 Memory Usage: 14kB
            -> Seq Scan on agencia (cost=0.00..5.60 rows=160 width=4) (actual time=0.012..0.041 rows=160
loops=6)
    Planning time: 0.283 ms
    Execution time: 193.116 ms
(13 registros)
```

```
cliente=> EXPLAIN ANALYZE SELECT COUNT(*) FROM documento JOIN agencia USING (agenciaID);
```

```
Aggregate (cost=156893.26..156893.27 rows=1 width=8) (actual time=1025.915..1025.915 rows=1 loops=1)
-> Hash Join (cost=7.60..151555.05 rows=2135284 width=0) (actual time=0.118..889.925 rows=2134578 loops=1)
    Hash Cond: (documento.agenciaID = agencia.agenciaID)
    -> Seq Scan on documento (cost=0.00..122195.84 rows=2135284 width=4) (actual time=0.038..409.942 rows=2135284
loops=1)
        -> Hash (cost=5.60..5.60 rows=160 width=4) (actual time=0.068..0.068 rows=160 loops=1)
            Buckets: 1024 Batches: 1 Memory Usage: 14kB
            -> Seq Scan on agencia (cost=0.00..5.60 rows=160 width=4) (actual time=0.008..0.040 rows=160 loops=1)
    Planning time: 0.337 ms
    Execution time: 1025.988 ms
```

```
Aggregate (cost=156893.26..156893.27 rows=1 width=8) (actual time=1020.201..1020.201 rows=1 loops=1)
-> Hash Join (cost=7.60..151555.05 rows=2135284 width=0) (actual time=0.109..883.553 rows=2134578 loops=1)
    Hash Cond: (documento.agenciaID = agencia.agenciaID)
    -> Seq Scan on documento (cost=0.00..122195.84 rows=2135284 width=4) (actual time=0.029..405.947 rows=2135284
loops=1)
        -> Hash (cost=5.60..5.60 rows=160 width=4) (actual time=0.067..0.067 rows=160 loops=1)
            Buckets: 1024 Batches: 1 Memory Usage: 14kB
            -> Seq Scan on agencia (cost=0.00..5.60 rows=160 width=4) (actual time=0.007..0.046 rows=160 loops=1)
    Planning time: 0.353 ms
    Execution time: 1020.273 ms
```

```
Aggregate (cost=156893.26..156893.27 rows=1 width=8) (actual time=1020.198..1020.198 rows=1 loops=1)
-> Hash Join (cost=7.60..151555.05 rows=2135284 width=0) (actual time=0.107..883.506 rows=2134578 loops=1)
    Hash Cond: (documento.agenciaID = agencia.agenciaID)
    -> Seq Scan on documento (cost=0.00..122195.84 rows=2135284 width=4) (actual time=0.031..405.432 rows=2135284
loops=1)
        -> Hash (cost=5.60..5.60 rows=160 width=4) (actual time=0.063..0.063 rows=160 loops=1)
            Buckets: 1024 Batches: 1 Memory Usage: 14kB
            -> Seq Scan on agencia (cost=0.00..5.60 rows=160 width=4) (actual time=0.007..0.045 rows=160 loops=1)
    Planning time: 0.345 ms
    Execution time: 1020.269 ms
```

- Parallel Bitmap Heap Scan
- Parallel Merge Join

- Parallel Hash Join
- Parallel Append
- CREATE TABLE AS
- CREATE MATERIALIZED VIEW
- SELECT INTO
- CREATE INDEX
  - Também em tabelas particionadas

- <https://blog.2ndquadrant.com/postgresql96-parallel-sequential-scan/>
- EDB – Query Parallelism In PostgreSQL – Dilip Kumar
- <http://rhaas.blogspot.com/2017/03/parallel-query-v2.html>
- PostgreSQL Documentation
- PostgreSQL 11 New Features – Hewlett Packard Enterprise



Rio Grande do Sul

GREAT  
PLACE  
TO  
WORK®

2017

Melhores Empresas  
Para Trabalhar™

BRASIL

AMANHÃ



# Obrigado!

54 3045-8100

[alvaro@atua.com.br](mailto:alvaro@atua.com.br)

[www.atua.com.br/trabalhe-conosco](http://www.atua.com.br/trabalhe-conosco)



[/atuasistemas](https://www.facebook.com/atuasistemas)