



Parallel Queries with PostgreSQL

Gianni Ciolli

PGDay.IT, 13 Ottobre 2017



Parallel Query

- The ability to use multiple CPUs in a single query
- Introduced in version 9.6
- Expanded in version 10

Parallel Feature Matrix

	9.6	10
Sequential Scan	✓	✓
B-tree Index Scan		✓
Bitmap Heap Scan		✓
Hash Join	✓	✓
Nested Loop Join	✓	✓
Merge Join		✓
Aggregation	✓	✓
Gather Merge		✓
More Parallel Queries in PL		✓



Single-CPU Query Example

- One int4 column
- 1M rows

Seq Scan on t

(cost=0.00..14425.00 rows=1000000 width=4)



Parallel Query Example

- One int4 column
- 1M rows

Gather

```
(cost=1000.00..10591.67 rows=1000000 width=4)
```

```
Workers Planned: 2
```

```
-> Parallel Seq Scan on t
```

```
(cost=0.00..8591.67 rows=416667 width=4)
```

The Gather Node

- There is one Gather node
- Separates single-CPU nodes from Parallel nodes




Figure: The Previous Example



Parallel Plan

- a.k.a. Partial Plan
- Executed by ≥ 1 parallel processes
- Each output row is generated exactly by one process
- Extra logic to split work properly across processes

Parallel Plan

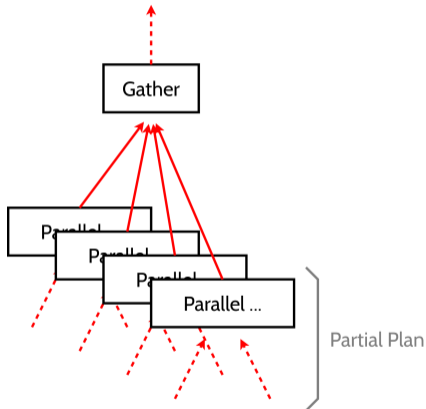


Figure: Generic Example



Parallel Safety

- Parallel Safe
 - Always OK
- Parallel Restricted
 - Only **above** a Gather node
- Parallel Unsafe
 - Never OK
 - Forces query to run on 1 CPU



How Many Parallel Workers?

-

$$1 + \log_3 \left\lfloor \frac{\text{(Size in blocks)}}{\text{min_parallel_*_size}} \right\rfloor$$



How Many Parallel Workers?

-

$$1 + \log_3 \left\lfloor \frac{\text{(Size in blocks)}}{\text{min_parallel_*_size}} \right\rfloor$$

???



How Many Parallel Workers?

-

$$1 + \log_3 \left\lfloor \frac{\text{(Size in blocks)}}{\text{min_parallel_*_size}} \right\rfloor$$

- Grows (slowly) with table size
- Grows by decreasing min_parallel_*_size
 - PostgreSQL 9.6:
 - min_parallel_relation_size
 - PostgreSQL 10:
 - min_parallel_table_scan_size
 - min_parallel_index_scan_size



Parallel Aggregation

- Done via Partial Aggregation
- Aggregate must:
 - be Parallel Safe
 - have a COMBINE method
 - State \rightarrow State \rightarrow State
- Each worker does a Partial Aggregation
- COMBINE aggregates partial results



And now...

Questions?



And then...

Thank you!

`gianni@2ndquadrant.com`
`@GianniCiolli`



Licence

This document is distributed under the **Creative Commons Attribution-Non commercial-ShareAlike 3.0 Unported** licence



A copy of the licence is available at the URL
<http://creativecommons.org/licenses/by-nc-sa/3.0/>

or you can write to

Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.