

# Перепланирование запросов



**Алена Рыбакина**

[a.rybakina@postgrespro.ru](mailto:a.rybakina@postgrespro.ru)

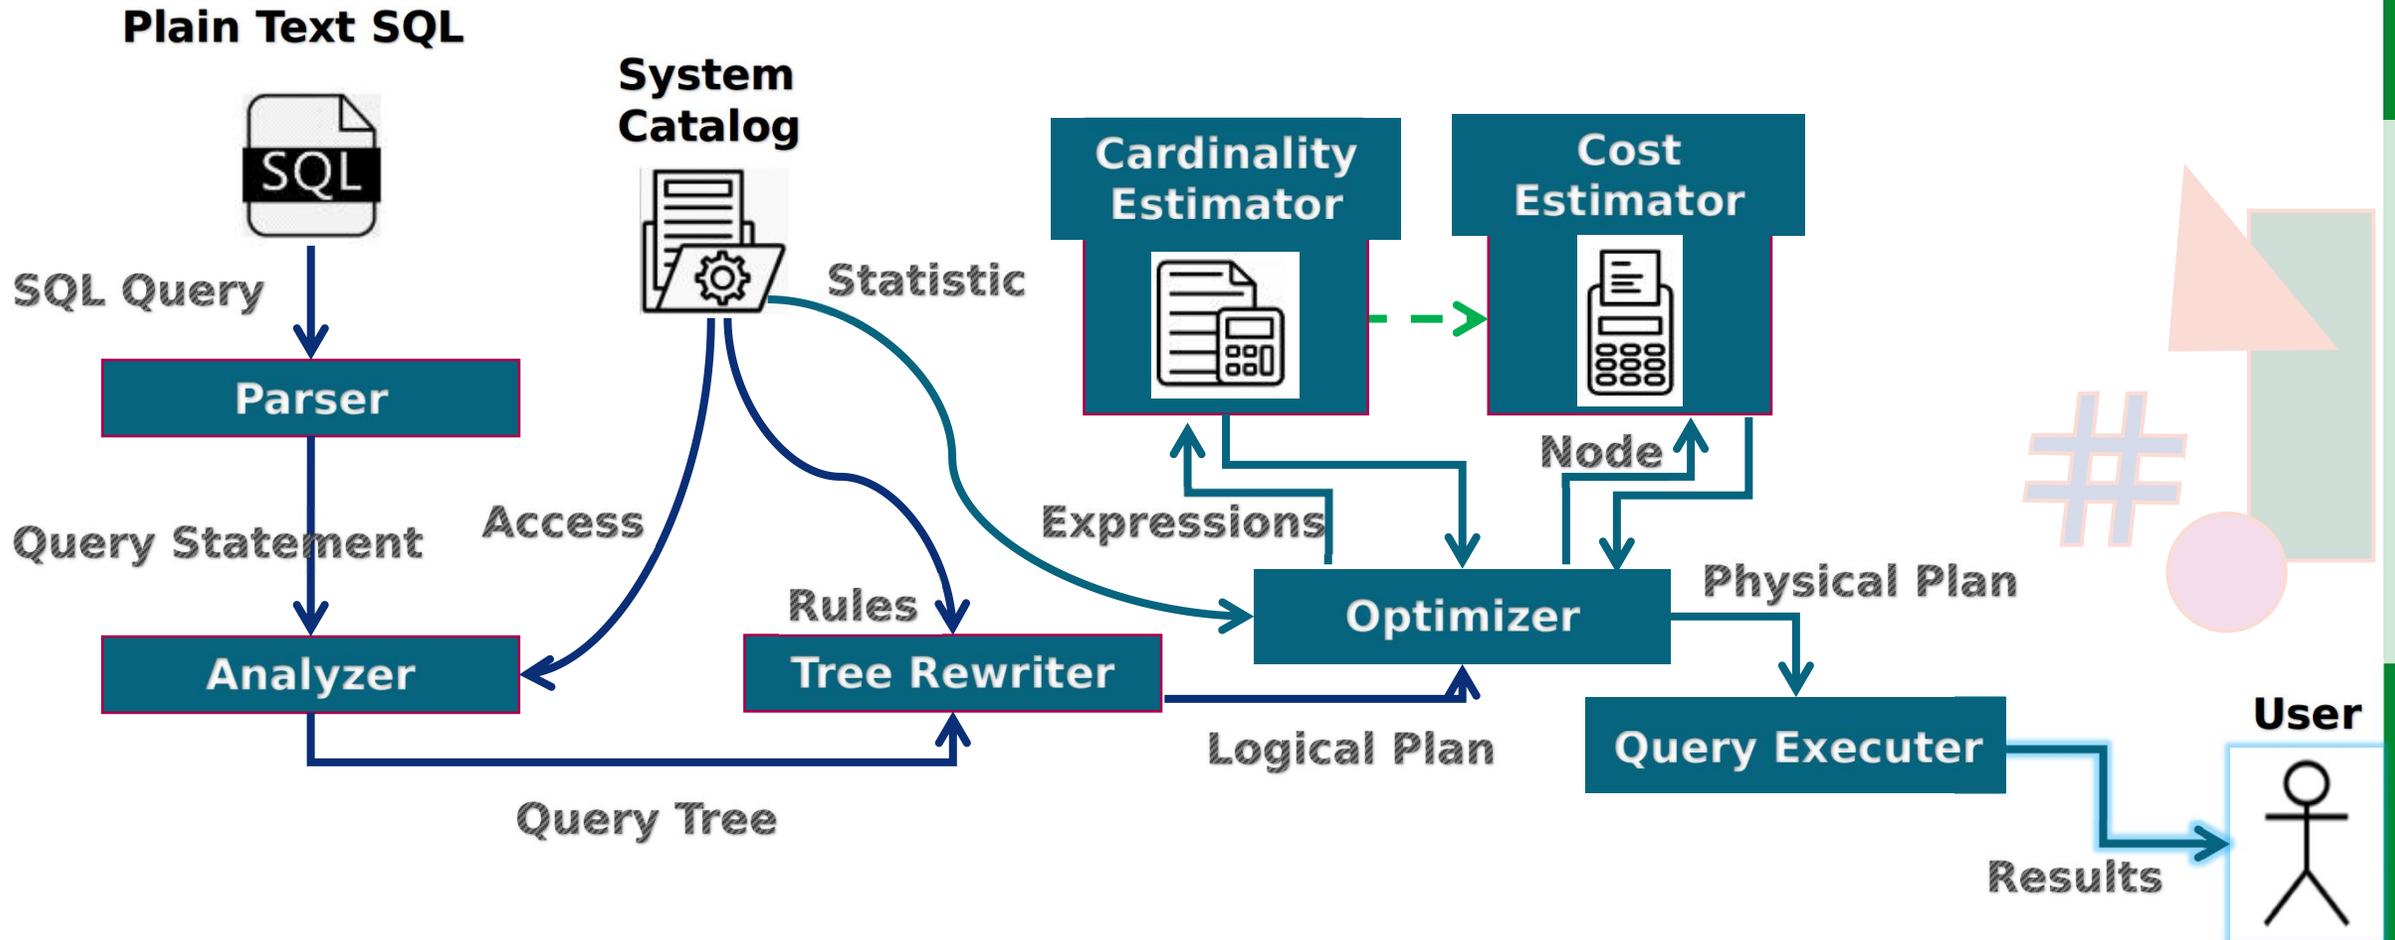
**Участники проекта:**

**Дамир Белялов, Андрей Лепихов, Алена Рыбакина**

- **Как работает оптимизатор**
- **Основная проблема оптимизатора**
- **Возможные решения**
- **Как работает Replan**
- **Результаты тестирования Replan**



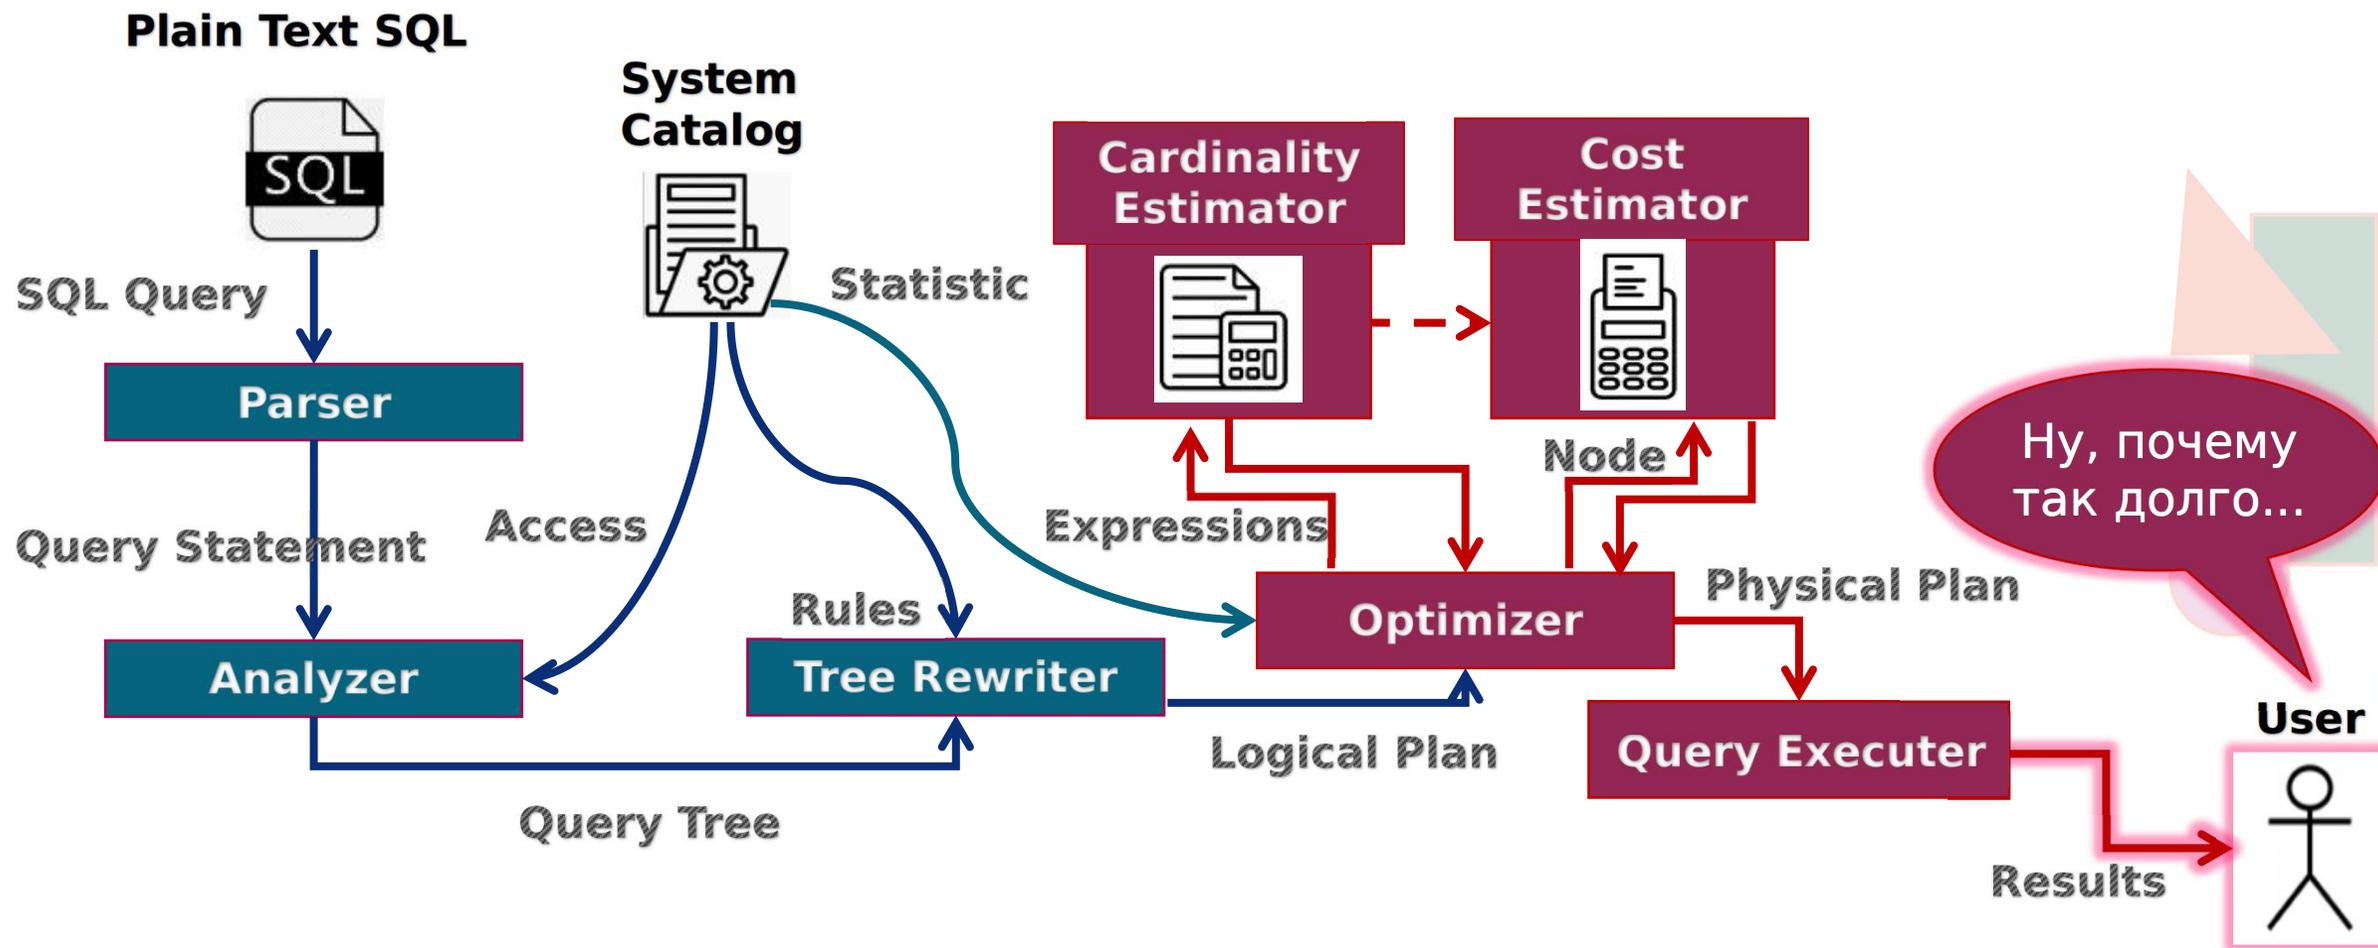
# Устройство оптимизатора



Как работает оптимизатор

Основная проблема оптимизатора

# Устройство оптимизатора



Как работает оптимизатор

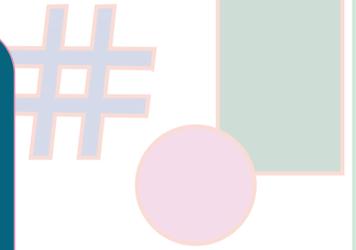
Основная проблема оптимизатора

## Причины неточной оценки кардинальности

использование  
устаревшей  
статистики

Предположение о  
равномерном  
распределении  
между столбцами

большое  
количество  
соединений в  
запросе



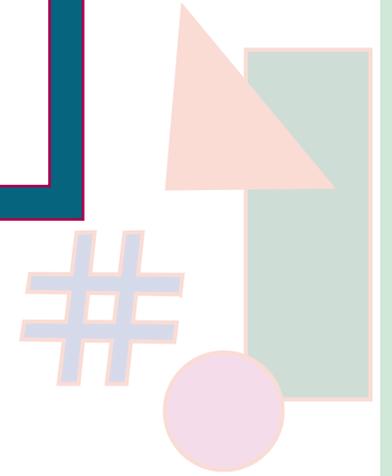
Основная проблема оптимизатора

Возможные решения

## Оптимизатор опирается на такие предположения

кардинальность  
предсказана всегда верно

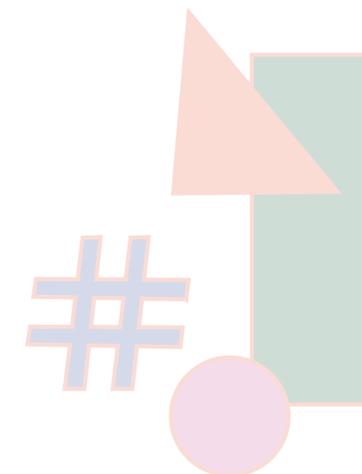
стоимость для плана  
всегда рассчитана  
корректно



Основная проблема оптимизатора

Возможные решения

Итак, проблема



# Решения от PostgreSQL

Указать функциональные зависимости

Указать статистику по выражениям

Указать число уникальных комбинаций значений в столбцах

Указать список наиболее частых комбинаций значений

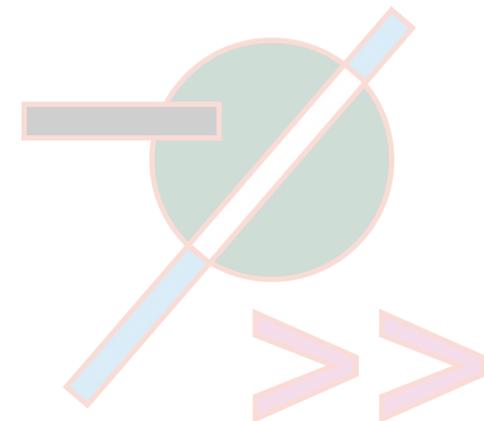


Расширенная статистика

Указать статистику индекса по выражению



Генетический алгоритм GEQO



Возможные решения

Как работает Replan

# О GEQO

Познакомимся с GEQO за 20 минут



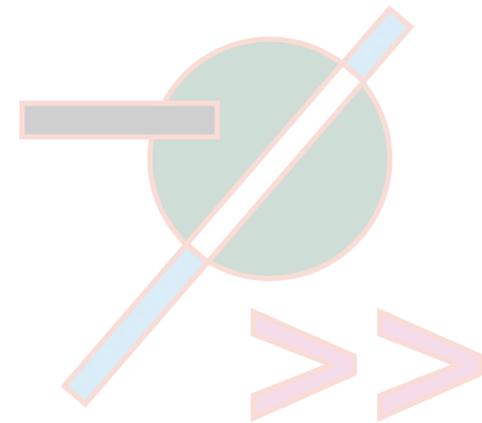
**ПАВЕЛ ТОЛМАЧЕВ**

Postgres Professional

Специалист образовательного отдела



**Генетический алгоритм GEQO**

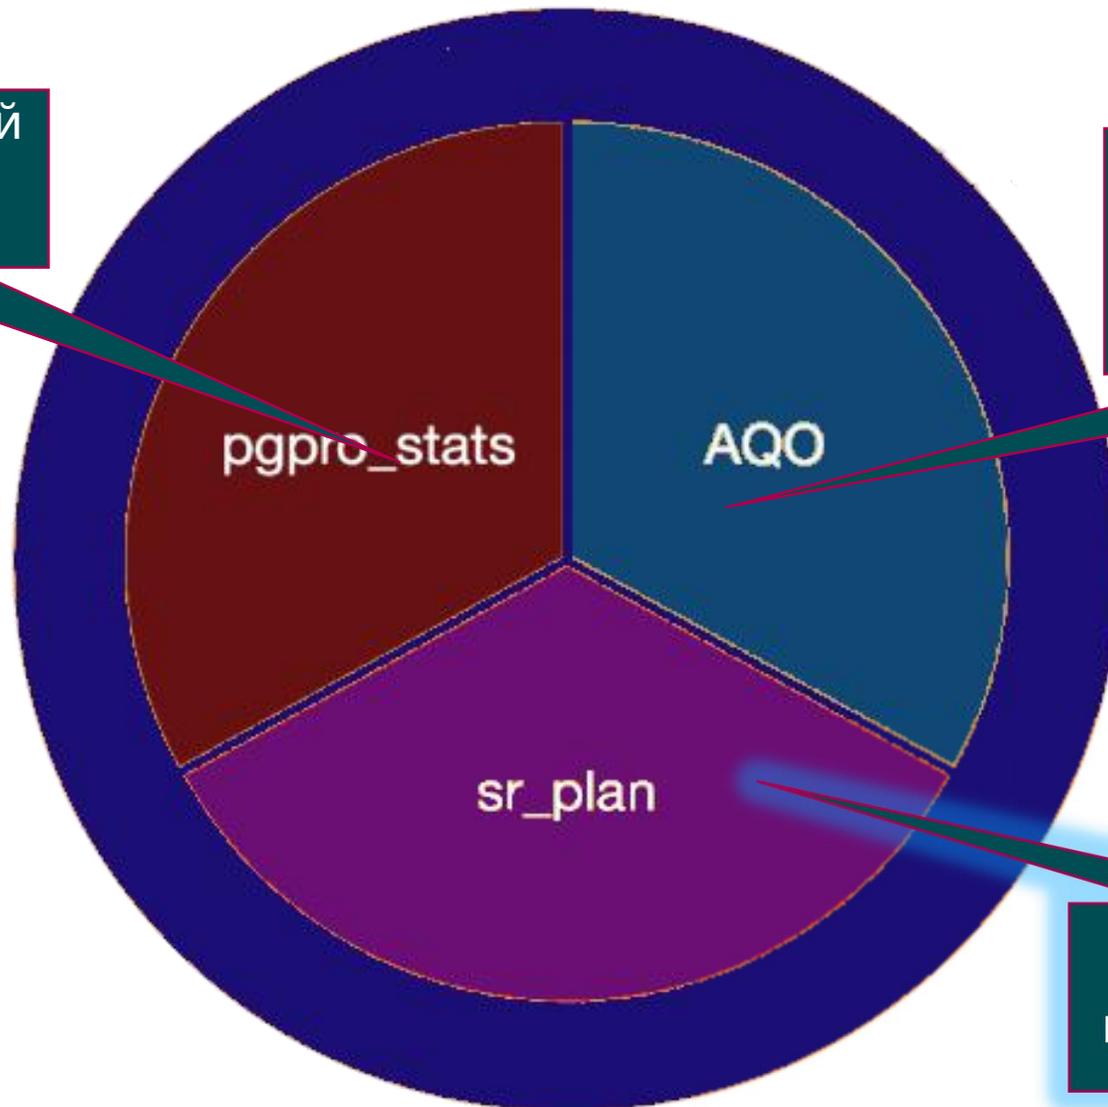


**Возможные решения**

**Как работает Replan**

# Решения от PostgresPro Enterprise

требуется ручной анализ и поиск лучшего плана



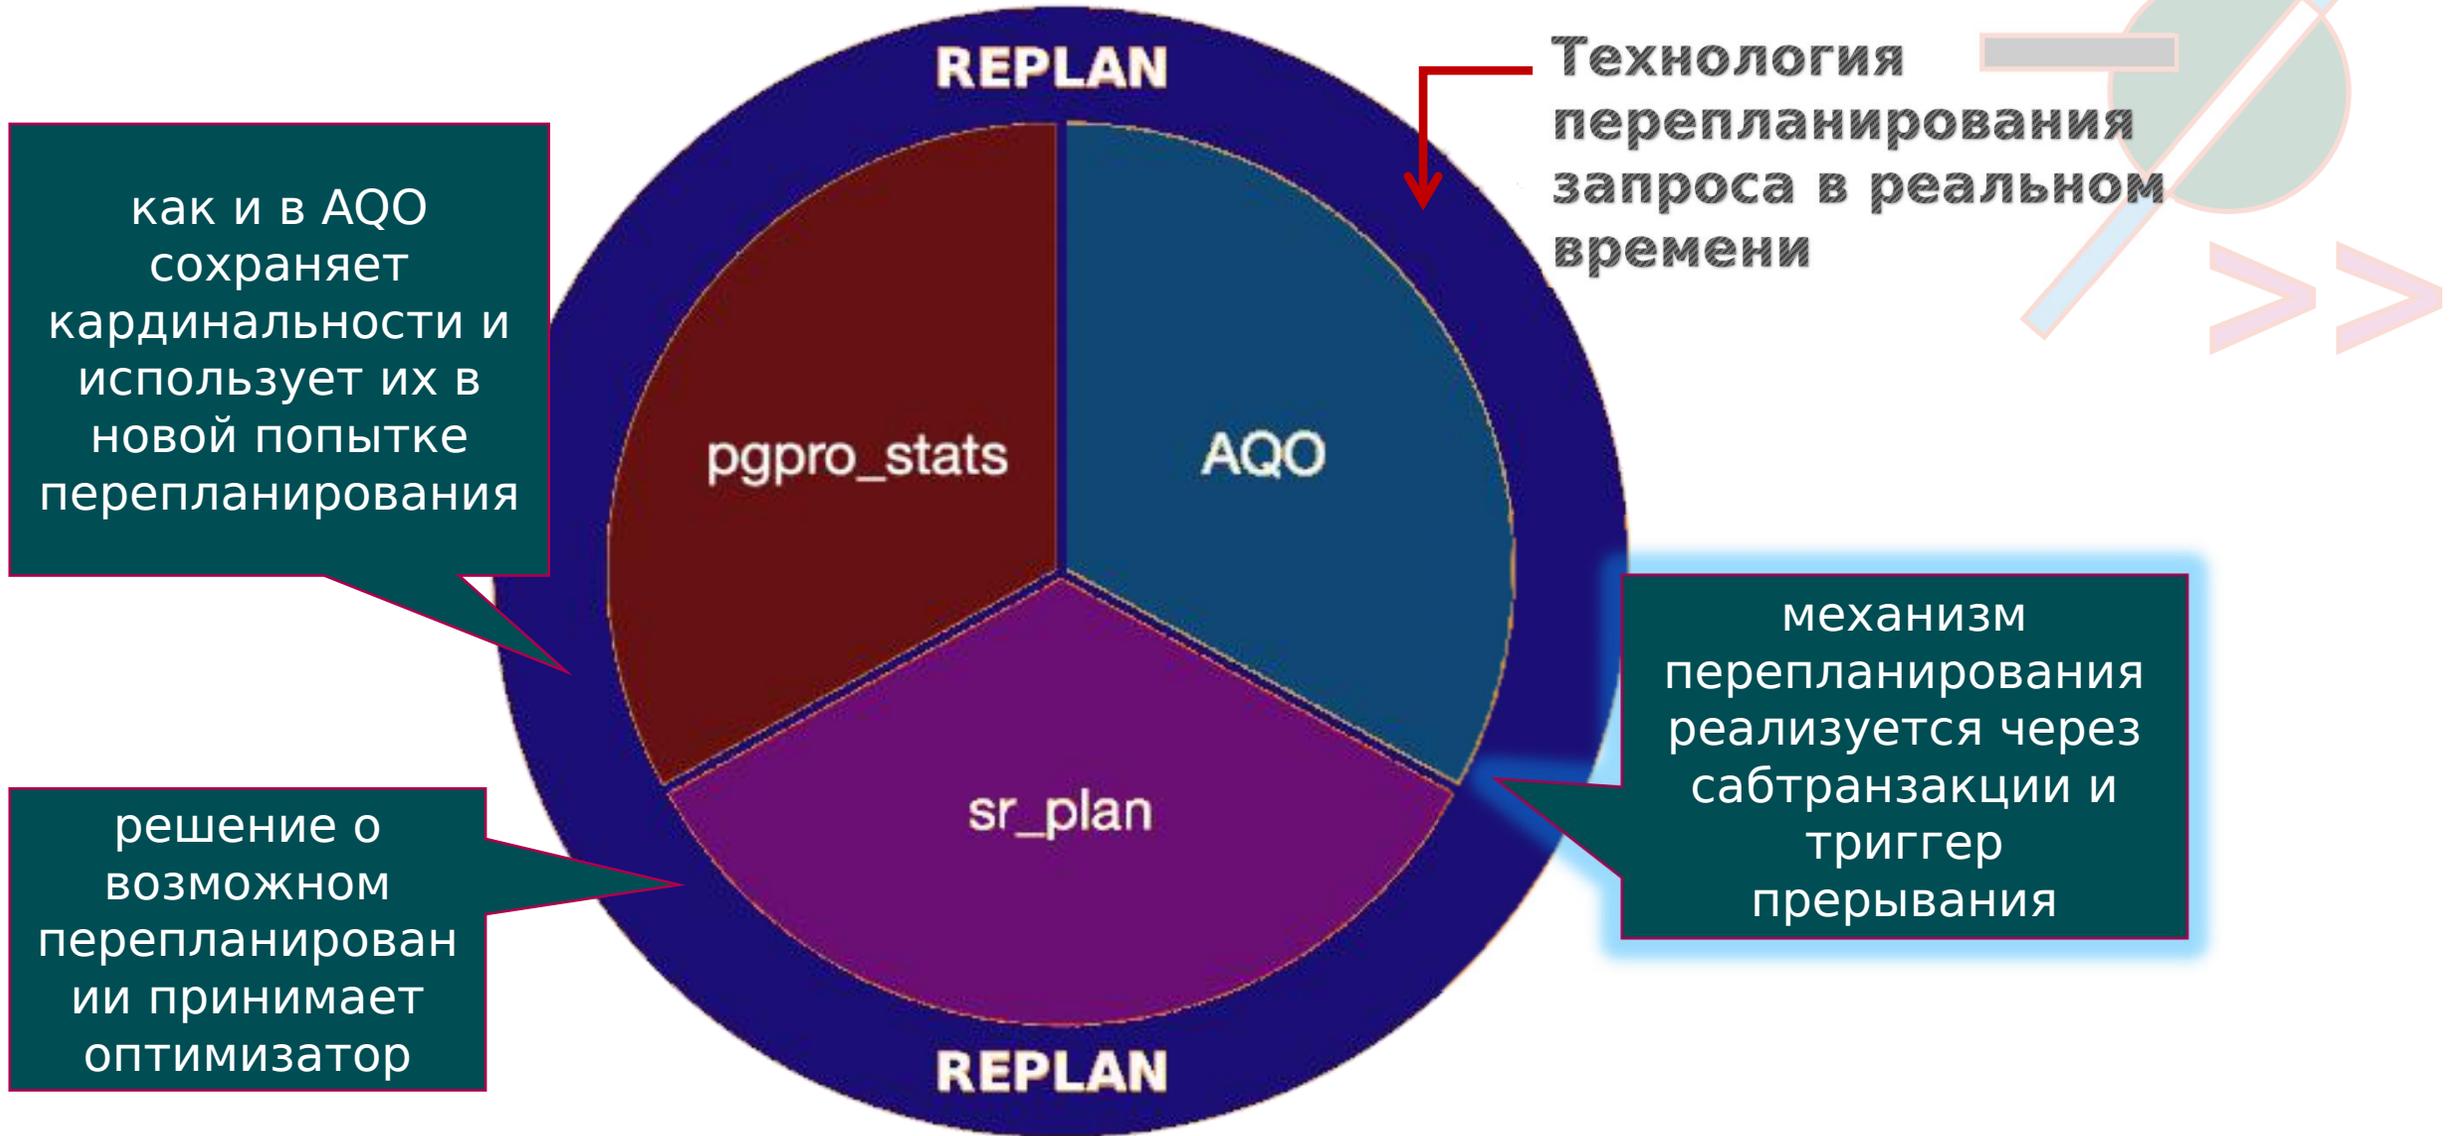
в основе него используется алгоритм Machine Learning, который может ошибаться

только фиксирует найденное решение

Возможные решения

Как работает Replan

# Решения от PostgresPro Enterprise



**Возможные решения**

**Как работает Replan**

# Две важные компоненты для реплана



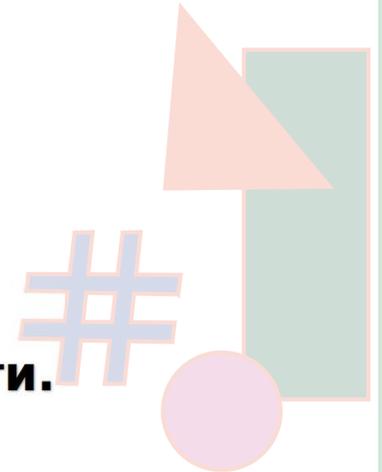
Replan  
Data

**Место хранения кардинальностей**



Trigger

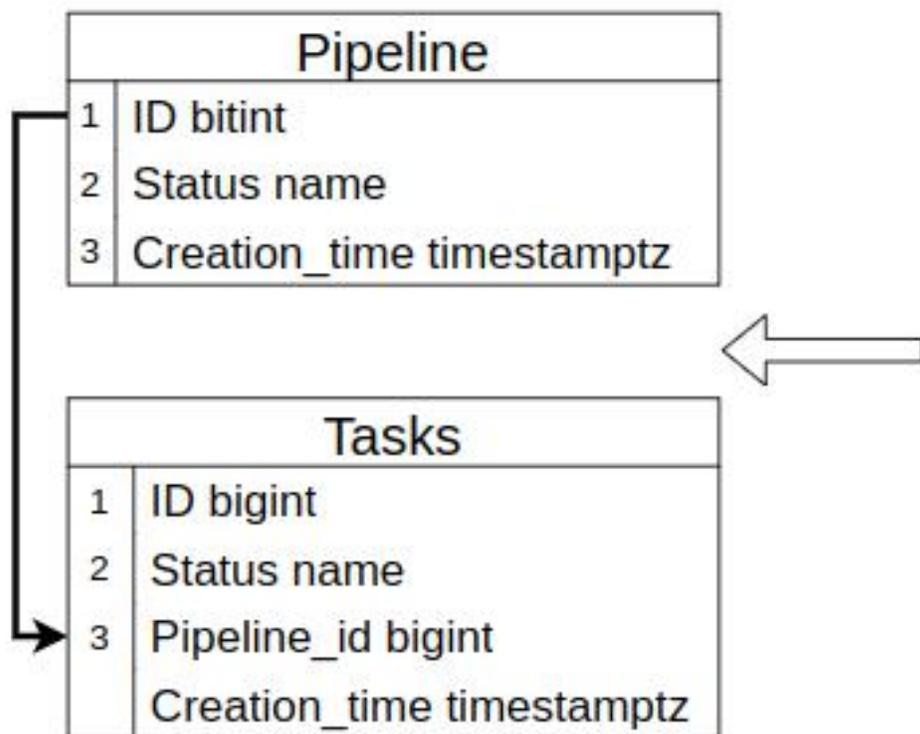
**Условие запуска перепланирования.  
Может быть по времени, по кардинальности, по памяти.  
Значения триггеров устанавливается пользователем.**



Как работает Replan

Результаты тестирования





- генерируется 50000 pipelines
- в каждом pipeline генерируется по 5 задач
- status pipeline может быть New или Started
- status задач может быть New или Completed

# Представление запроса

```
EXPLAIN ANALYZE SELECT count(1) FROM pipeline p, tasks t
WHERE t.status = 'NEW' AND p.id = t.pipeline_id AND t.creation_time > now() - interval '1 hour';
```

---

```
Aggregate (rows=1) (rows=1)
-> Nested Loop (rows=1) (rows=629973)
    -> Index Scan using tasks_status_creation_time_idx1 on tasks t (rows=1) (rows=1889920)
        Index Cond: ((status = 'NEW'::text) AND (creation_time > (now() -
'01:00:00'::interval)))
    -> Index Only Scan using pipeline_pkey on pipeline p (rows=1) (rows=629973)
        Index Cond: (id = t.pipeline_id)
```

---

Planning Time: 1.045 ms

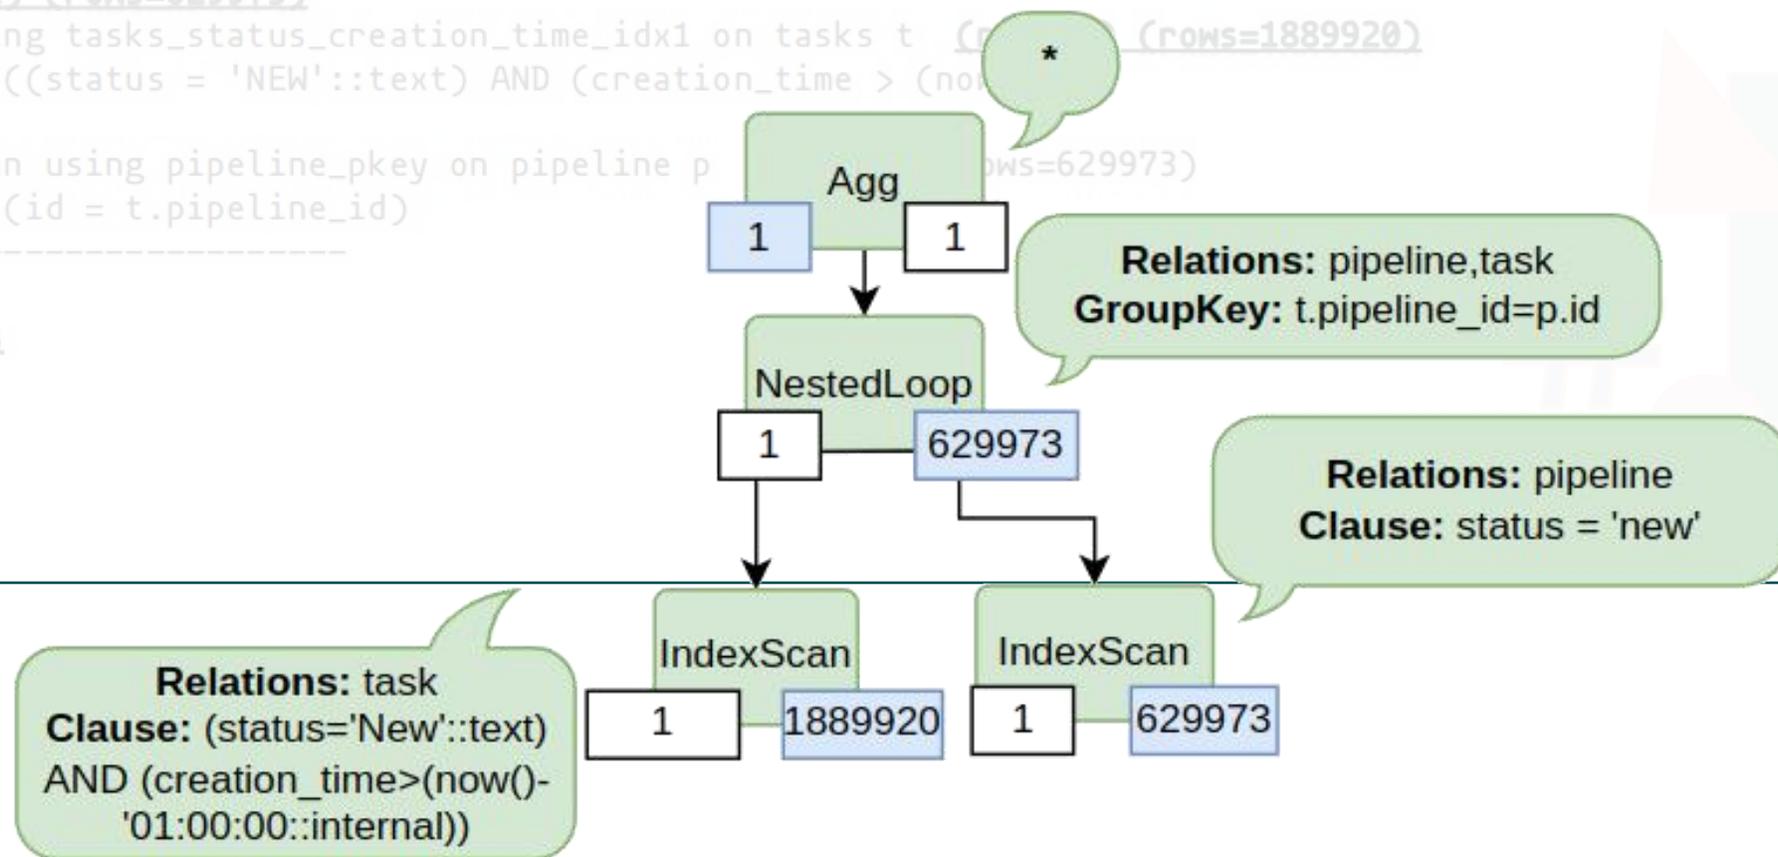
**Execution Time: 9595.701 ms**

# Представление запроса

```
EXPLAIN ANALYZE SELECT count(1) FROM pipeline p, tasks t
WHERE t.status = 'NEW' AND p.id = t.pipeline_id AND t.creation_time > now() - interval '1 hour';
```

```
Aggregate (rows=1) (rows=1)
-> Nested Loop (rows=1) (rows=629973)
  -> Index Scan using tasks_status_creation_time_idx1 on tasks t (rows=1889920)
      Index Cond: ((status = 'NEW'::text) AND (creation_time > (now() - interval '01:00:00'::interval)))
  -> Index Only Scan using pipeline_pkey on pipeline p (rows=629973)
      Index Cond: (id = t.pipeline_id)
```

```
Planning Time: 1.045 ms
Execution Time: 9595.701 ms
```



Как работает Replan

Результаты тестирования

# Explain запроса с репланом в терминале

```
EXPLAIN ANALYZE SELECT count(1) FROM pipeline p, tasks t
WHERE t.status = 'NEW' AND p.id = t.pipeline_id AND t.creation_time > now() - interval '1 hour';
```

---

```
Finalize Aggregate (rows=1) (rows=1)
-> Gather (rows=2) (rows=3)
    -> Partial Aggregate (rows=1) (rows=1)
        -> Hash Join (rows=8355) (rows=629973)
            Hash Cond: (p.id = t.pipeline_id)
                -> Parallel Seq Scan on pipeline p (rows=764484) (rows=628661)
                -> Hash (rows=240357) (rows=1889920)
                    -> Index Scan using tasks_status_creation_time_idx1 on tasks t
                                                                (rows=240357) (rows=1889920)
                        Index Cond: ((status = 'NEW'::text) AND (creation_time > (now() -
'01:00:00'::interval)))
```

---

```
Planning Time: 0.452 ms
Replanning Attempts: 5
```

```
Execution Time: 2229.972 ms
Total Execution Time: 2964.835 ms
```

# Лог реплана

2023-09-21 13:31:49.469 MSK [8320] LOG: Replanning triggered by timeout 100 (0-th shift) ms.

Attempt: 1.

Duration: 158.919477 ms plan:

Query Text: explain (analyze) select count(1) from pipeline p, tasks t...

Finalize Aggregate (cost=30098.31..30098.32 rows=1 width=8) (actual time=162.220..162.220 rows=0 loops=1) (early terminated)

NodeSign: 4304986849338732187

Cardinality: -1

Groups Number: -1

Output: count(1)...

2023-09-21 13:31:49.957 MSK [8320] LOG: Replanning triggered by timeout 100 (0-th shift) ms.

Attempt: 4.

Duration: 156.986469 ms plan:

Query Text: explain (analyze) select count(1) from pipeline p, tasks t...

Finalize Aggregate (cost=41139.66..41139.67 rows=1 width=8) (actual time=160.432..160.432 rows=0 loops=1) (early terminated)

NodeSign: 4304986849338732187

Cardinality: -1

Groups Number: -1

Output: count(1)...

Как работает Replan

Результаты тестирования

```
EXPLAIN ANALYZE SELECT count(1) FROM pipeline p, tasks t
  WHERE p.status = 'STARTED' AND p.id = t.pipeline_id and t.status IN ('NEW', 'COMPLETED')
  AND t.creation_time > now() - interval '1 hour';
```

---

```
Aggregate  (rows=1) (rows=1)
  -> Nested Loop  (rows=9) (rows=25000)
    -> Index Scan using pipeline_status_idx on pipeline p  (rows=1) (rows=5000)
        Index Cond: (status = 'STARTED'::text)
    -> Index Scan using tasks_pipeline_id_idx on tasks t  (rows=5) (rows=5)
        Index Cond: (pipeline_id = p.id)
        Filter: ((status = ANY ('{NEW,COMPLETED}'::text[])) AND (creation_time > (now() -
'01:00:00'::interval)))
```

---

```
Planning Time: 3.043 ms
Execution Time: 43.531 ms
(9 rows)
```

```
EXPLAIN ANALYZE SELECT count(1) FROM pipeline p, tasks t
  WHERE p.status = 'STARTED' AND p.id = t.pipeline_id and t.status IN ('NEW', 'COMPLETED')
  AND t.creation_time > now() - interval '1 hour';
```

```
-----
Aggregate  (rows=1) (rows=1)
  -> Hash Join  (rows=491) (rows=25000)
      Hash Cond: (t.pipeline_id = p.id)
        -> Index Scan using tasks_status_creation_time_idx on tasks t  (rows=24752) (rows=25000)
            Index Cond: ((status = 'NEW'::text) AND (creation_time > (now() -
'01:00:00'::interval)))
          -> Hash  (rows=99) (rows=5000)
              -> Index Scan using pipeline_status_idx on pipeline p  (rows=99) (rows=5000)
                  Index Cond: (status = 'STARTED'::text)
```

Planning Time: 0.414 ms

**Execution Time: 22.380 ms**

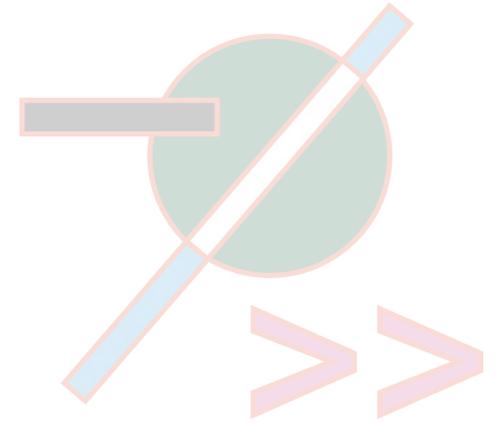
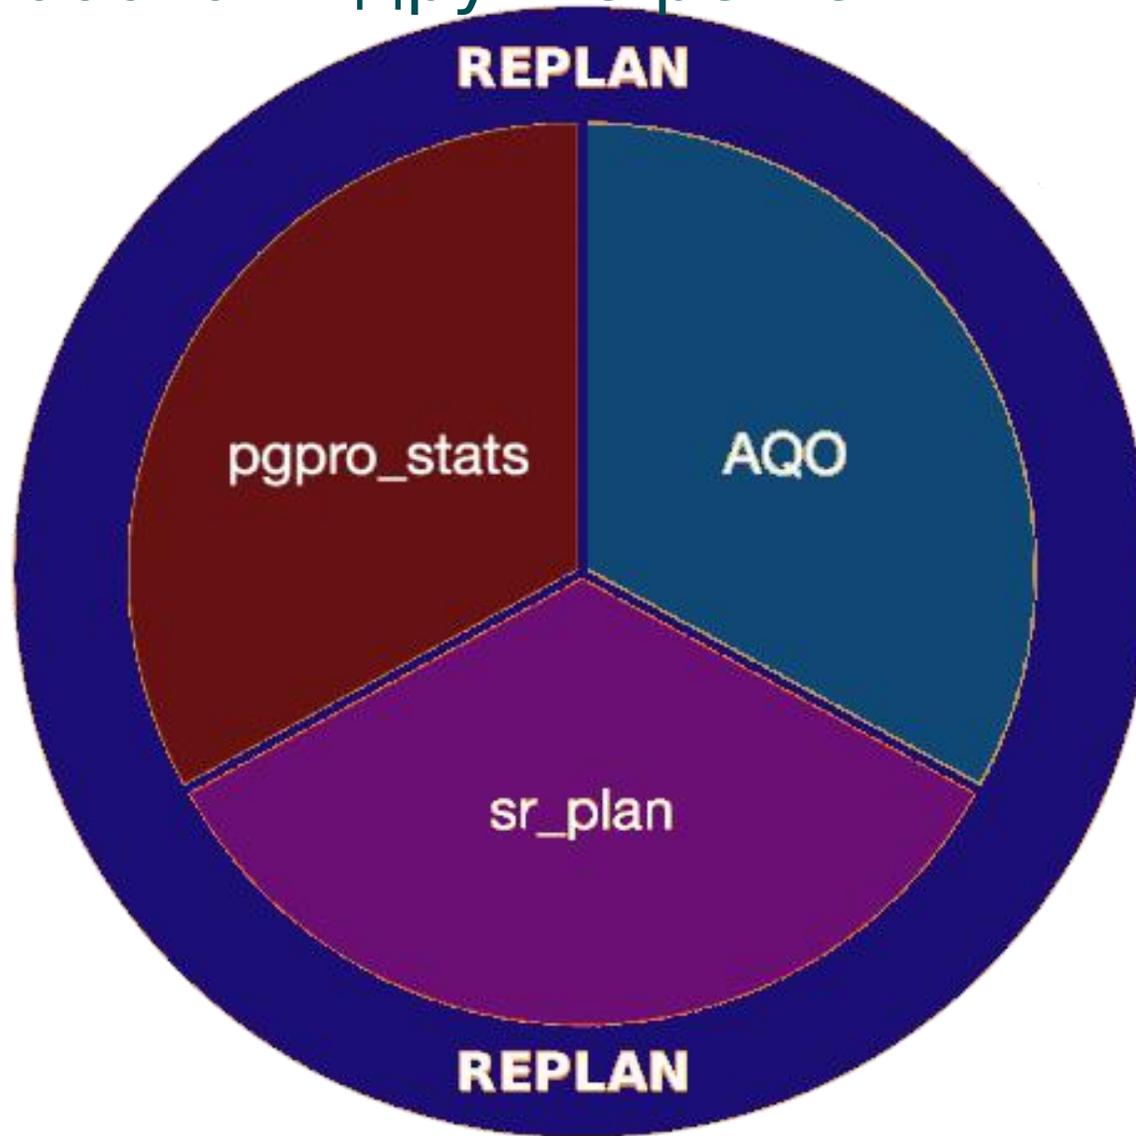
Replanning Attempts: 1

**Total Execution Time: 1025.460 ms**

**1. Запрос просто может быть написан плохо**

**2. Replan потребует большое количество итераций, чтобы найти лучший план**

Стоит попробовать другие решения

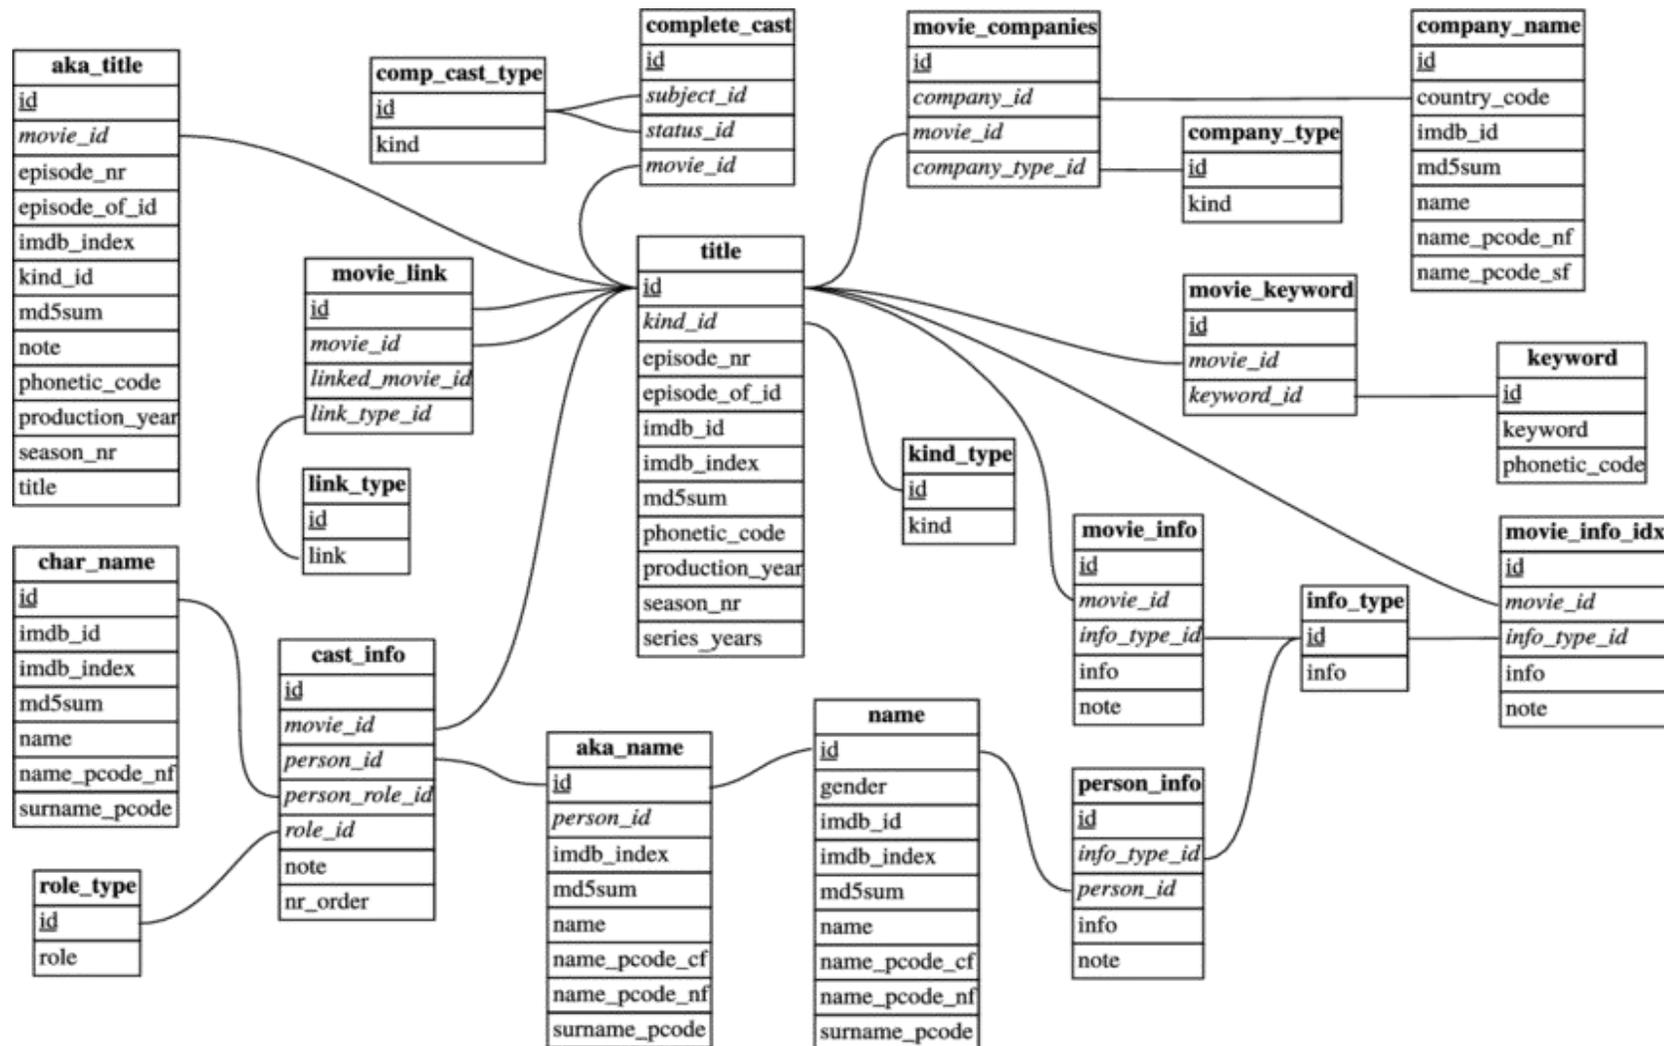


Как работает Replan

Результаты тестирования

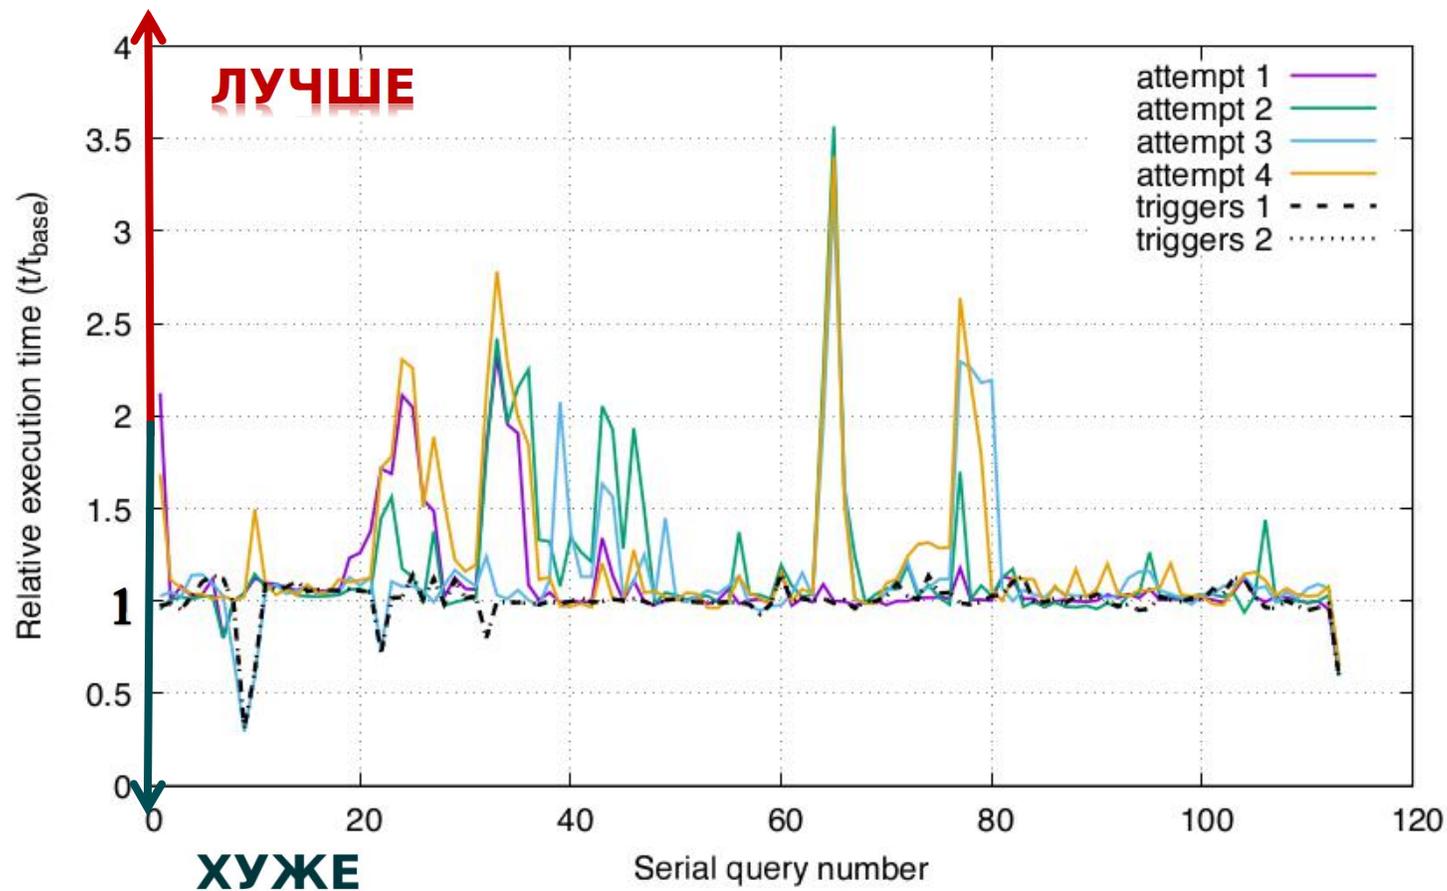
## Join Order Benchmark (JOB)

- состоит из 113 запросов по IMDB
- каждый запрос может иметь от 3 до 16 соединений
- запросы отвечают на логичные вопросы любителя кино
- для оптимизатора запросы сложные из-за большого количества объединений и корреляций



## Время выполнения запросов

Время выполнения запроса в JOB может сильно меняться от запуска к запуску

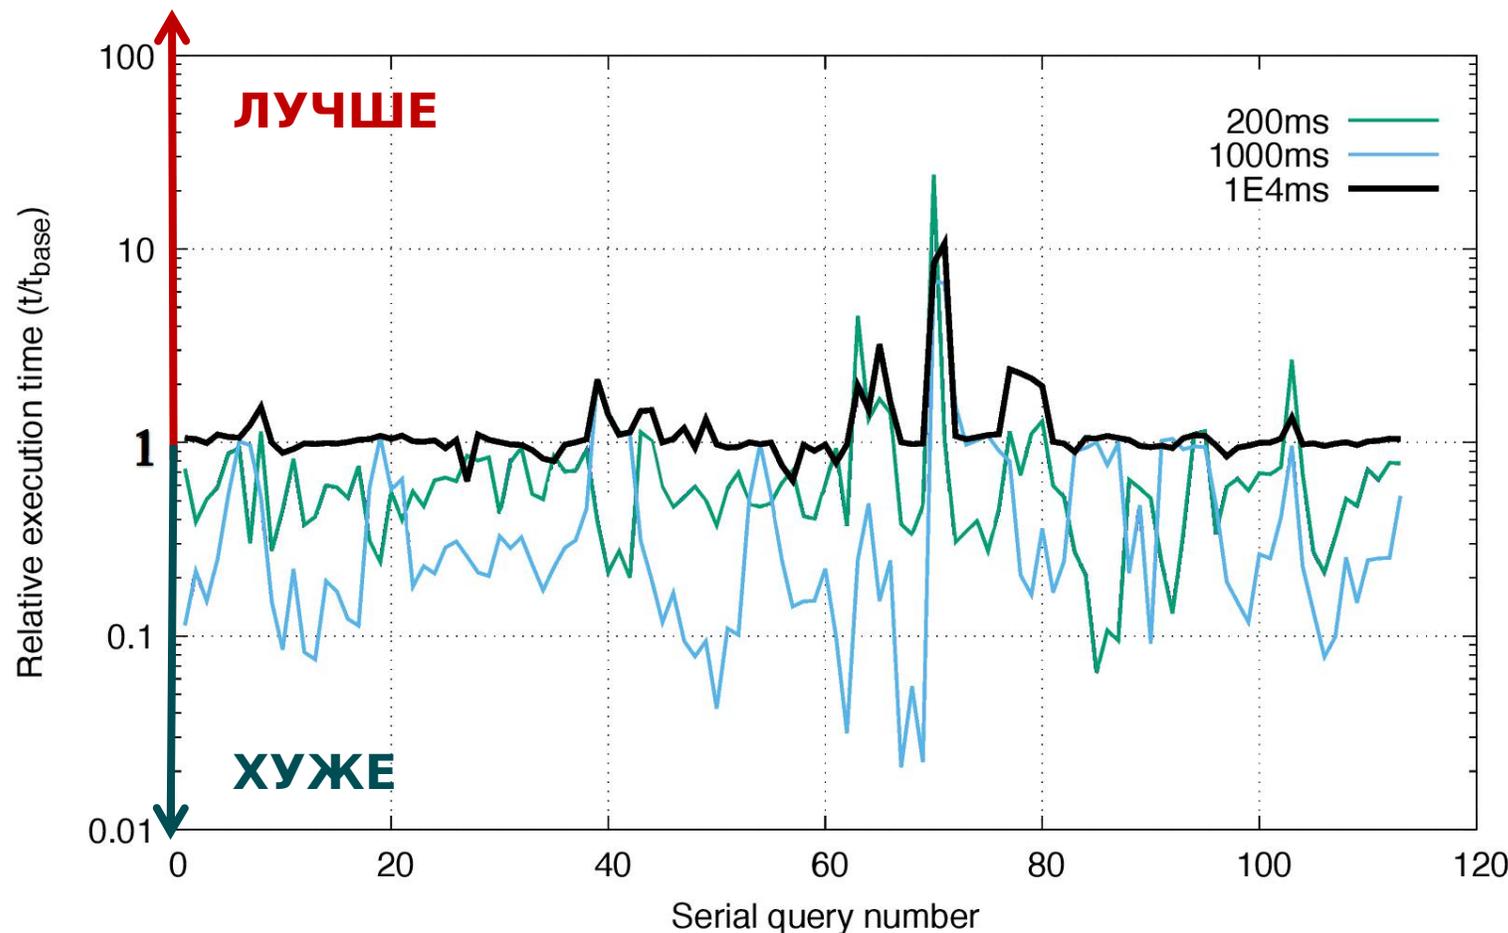


Результаты тестирования

# Тестирование на Join Order Benchmark

Время выполнения запросов

Время выполнения запроса после перепланирования может ухудшиться

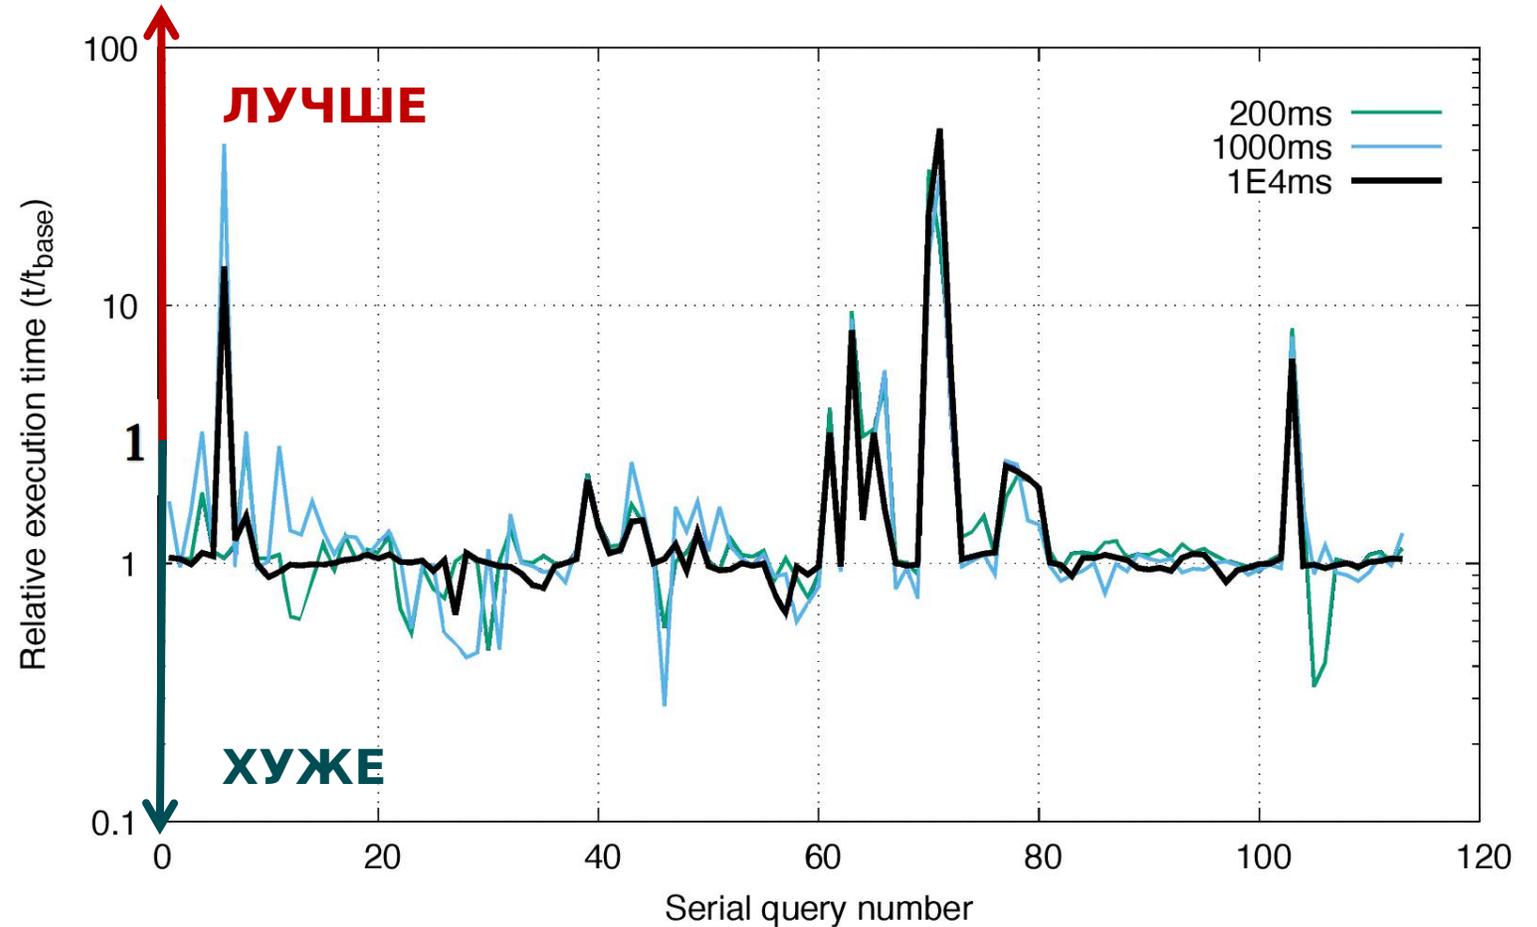


Результаты тестирования

# Тестирование на Join Order Benchmark

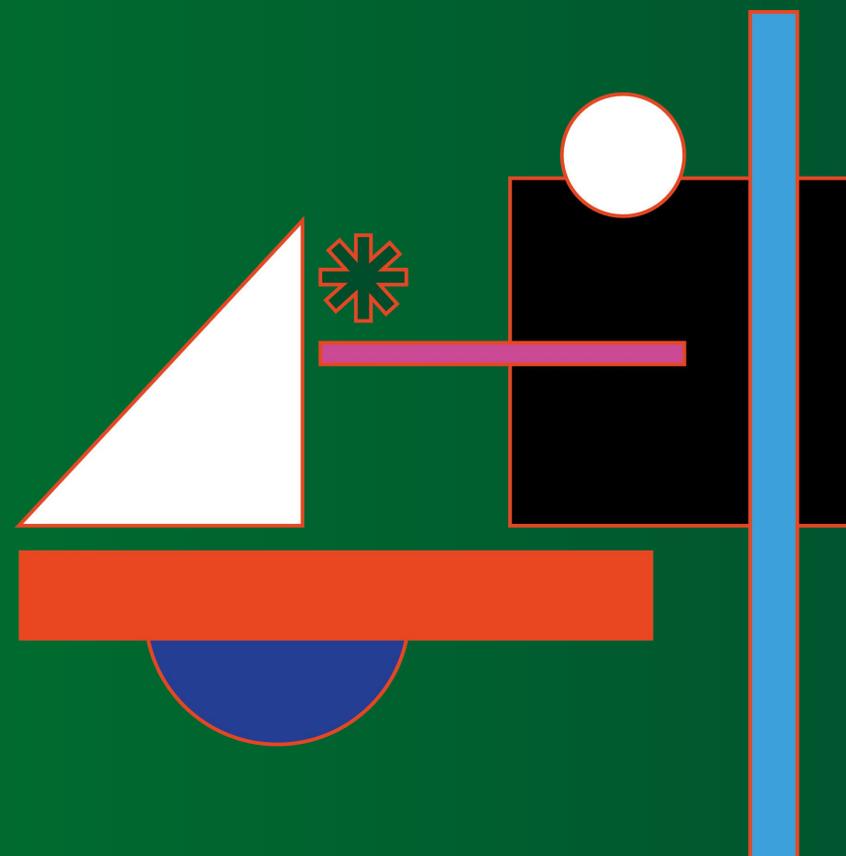
Итоговое время потраченное на перепланирование запросов

Перепланирование –  
крайняя мера



Результаты тестирования

# А теперь ДЕМО



```
(base) alena on alena-ThinkPad-L13-Gen-2 in ~
$ tmp_install/bin/psql -d alena
psql (12.15, server 17devel)
WARNING: psql major version 12, server major version 17.
        Some psql features might not work.
Type "help" for help.
```

```
alena=# DROP TABLE IF EXISTS a,b CASCADE;
CREATE TABLE a WITH (autovacuum_enabled = false) AS
  SELECT gs AS x FROM generate_series(1,1E5) AS gs;
CREATE TABLE b WITH (autovacuum_enabled = false) AS
  SELECT gs AS x FROM generate_series(1,1E5) AS gs;
DROP TABLE
SELECT 100000
SELECT 100000
alena=# █
```

(base) alena on alena-ThinkPad-L13-Gen-2 in ~

```
$ tmp_install/bin/psql -d alena
```

```
psql (12.15, server 17devel)
```

```
WARNING: psql major version 12, server major version 17.
```

```
Some psql features might not work.
```

```
Type "help" for help.
```

```
alena=# DROP TABLE IF EXISTS a,b CASCADE;
```

```
CREATE TABLE a WITH (autovacuum_enabled = false) AS
```

```
SELECT gs AS x FROM generate_series(1,1E5) AS gs;
```

```
CREATE TABLE b WITH (autovacuum_enabled = false) AS
```

```
SELECT gs AS x FROM generate_series(1,1E5) AS gs;
```

```
DROP TABLE
```

```
SELECT 100000
```

```
SELECT 100000
```

```
alena=# ANALYZE a,b;
```

```
INSERT INTO a SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
```

```
INSERT INTO b SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
```

```
ANALYZE
```

```
INSERT 0 100
```

```
INSERT 0 100
```

```
alena=#
```

```
psql (12.15, server 17devel)
```

```
WARNING: psql major version 12, server major version 17.
```

```
Some psql features might not work.
```

```
Type "help" for help.
```

```
alena=# DROP TABLE IF EXISTS a,b CASCADE;
```

```
CREATE TABLE a WITH (autovacuum_enabled = false) AS
```

```
  SELECT gs AS x FROM generate_series(1,1E5) AS gs;
```

```
CREATE TABLE b WITH (autovacuum_enabled = false) AS
```

```
  SELECT gs AS x FROM generate_series(1,1E5) AS gs;
```

```
DROP TABLE
```

```
SELECT 100000
```

```
SELECT 100000
```

```
alena=# ANALYZE a,b;
```

```
INSERT INTO a SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
```

```
INSERT INTO b SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
```

```
ANALYZE
```

```
INSERT 0 100
```

```
INSERT 0 100
```

```
alena=# SET show_node_sign = 'off';
```

```
SET query_inadequate_execution_time = -1;
```

```
SET replan_regression_mode = 'off';
```

```
SET
```

```
SET
```

```
SET
```

```
alena=#
```

Type "help" for help.

```
alena=# DROP TABLE IF EXISTS a,b CASCADE;
CREATE TABLE a WITH (autovacuum_enabled = false) AS
  SELECT gs AS x FROM generate_series(1,1E5) AS gs;
CREATE TABLE b WITH (autovacuum_enabled = false) AS
  SELECT gs AS x FROM generate_series(1,1E5) AS gs;
DROP TABLE
SELECT 100000
SELECT 100000
alena=# ANALYZE a,b;
INSERT INTO a SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
INSERT INTO b SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
ANALYZE
INSERT 0 100
INSERT 0 100
alena=# SET show_node_sign = 'off';
SET query_inadequate_execution_time = -1;
SET replan_regression_mode = 'off';
SET
SET
SET
alena=# EXPLAIN (ANALYZE, COSTS ON)
SELECT count(*) FROM a,b WHERE a.x=b.x AND b.x = 1
; -- Bad plan and execution time
```

## QUERY PLAN

```
-----  
Aggregate (cost=3396.01..3396.02 rows=1 width=8) (actual time=2453.918..2453.919 rows=1 loops=1)  
-> Nested Loop (cost=0.00..3396.01 rows=1 width=0) (actual time=0.065..2452.964 rows=10201 loops=1)  
    I-> Seq Scan on a (cost=0.00..1698.00 rows=1 width=6) (actual time=0.037..27.978 rows=101 loops=1)  
        Filter: (x = '1'::numeric)  
        Rows Removed by Filter: 99999  
    -> Seq Scan on b (cost=0.00..1698.00 rows=1 width=6) (actual time=0.007..23.998 rows=101 loops=101)  
        Filter: (x = '1'::numeric)  
        Rows Removed by Filter: 99999  
Planning Time: 0.514 ms  
Execution Time: 2453.987 ms  
(10 rows)
```

**(END)**

```
SELECT gs AS x FROM generate_series(1,1E5) AS gs;
DROP TABLE
SELECT 100000
SELECT 100000
alena=# ANALYZE a,b;
INSERT INTO a SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
INSERT INTO b SELECT 1 AS x FROM generate_series(1,1E2) AS gs;
ANALYZE
INSERT 0 100
INSERT 0 100
alena=# SET show_node_sign = 'off';
SET query_inadequate_execution_time = -1;
SET replan_regression_mode = 'off';
SET
SET
SET
alena=# EXPLAIN (ANALYZE, COSTS ON)
SELECT count(*) FROM a,b WHERE a.x=b.x AND b.x = 1
; -- Bad plan and execution time
alena=# SET show_node_sign = 'on';
SET query_inadequate_execution_time = 500;
SET
SET
alena=# EXPLAIN (ANALYZE, COSTS ON)
SELECT count(*) FROM a,b WHERE a.x=b.x AND b.x = 1
; -- Better plan and execution time
```

## QUERY PLAN

```
-----  
-----  
Aggregate (cost=3424.85..3424.86 rows=1 width=8) (actual time=47.884..47.885 rows=1 loops=1)  
  NodeSign: 17722406278722554782  
    -> Nested Loop (cost=0.00..3420.05 rows=1921 width=0) (actual time=0.025..47.208 rows=10201 loops=1)  
      NodeSign: 17684322539419885852  
        -> Seq Scan on b (cost=0.00..1698.00 rows=96 width=6) (actual time=0.013..22.580 rows=101 loops=1)  
          NodeSign: 170651357001224188  
          Filter: (x = '1'::numeric)  
          Rows Removed by Filter: 99999  
        -> Materialize (cost=0.00..1698.10 rows=20 width=6) (actual time=0.000..0.235 rows=101 loops=101)  
          NodeSign: subordinate  
          -> Seq Scan on a (cost=0.00..1698.00 rows=20 width=6) (actual time=0.009..23.242 rows=101 loops=1)  
            NodeSign: 7777922676653275221  
            Filter: (x = '1'::numeric)  
            Rows Removed by Filter: 99999
```

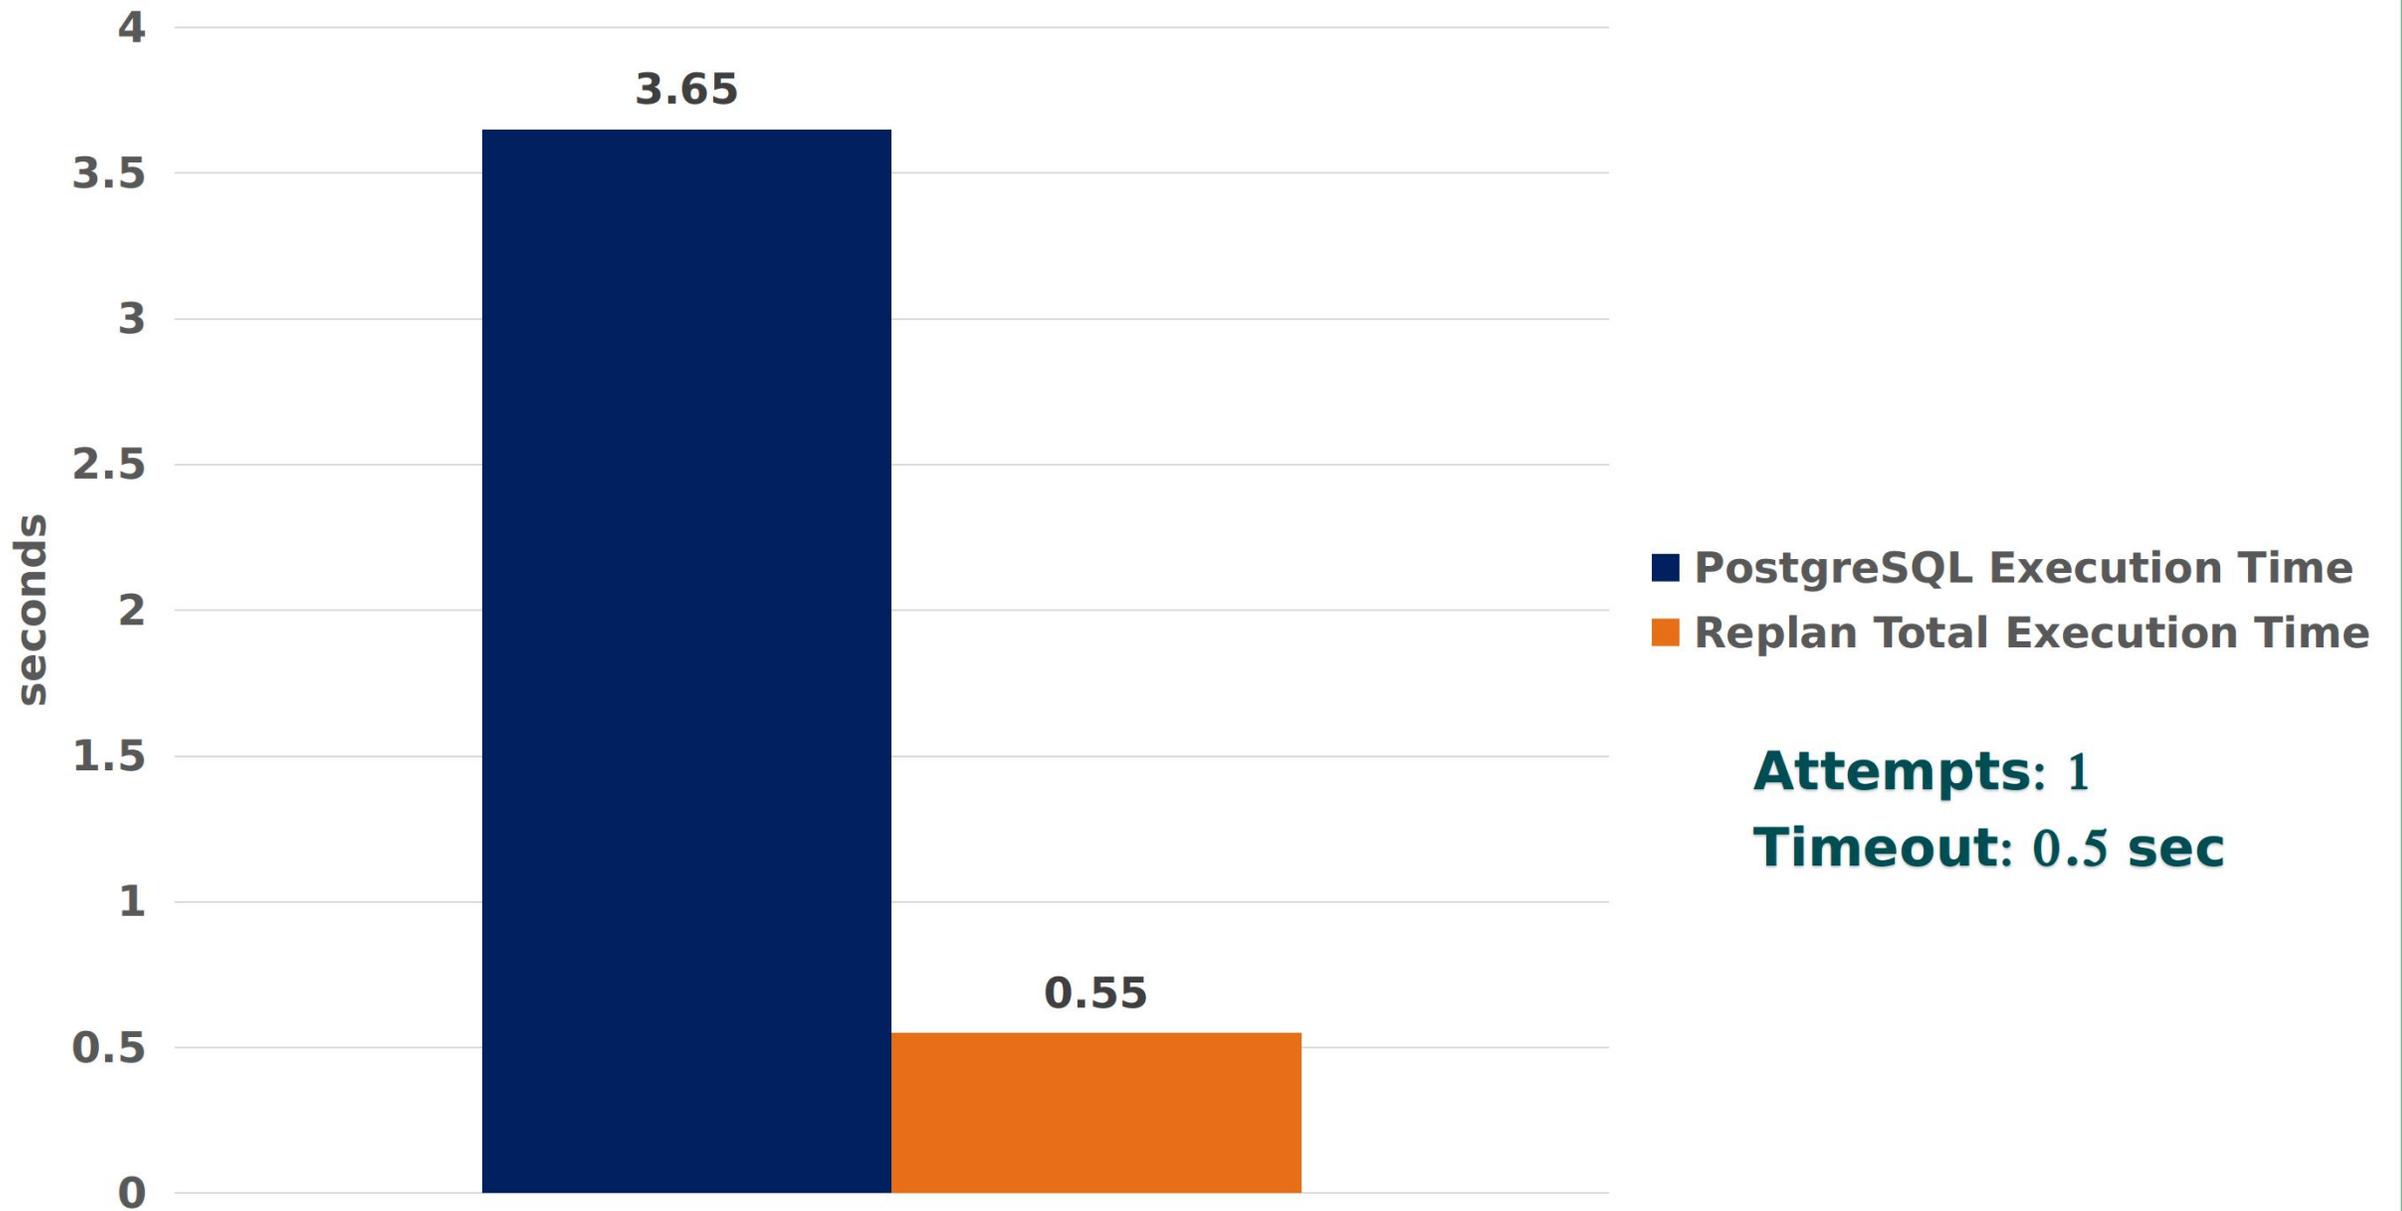
Planning Time: 0.146 ms

Execution Time: 47.942 ms

Replan Active: true

:█

# Query execution time



```
FROM company_name AS cn,  
     company_type AS ct,  
     keyword AS k,  
     link_type AS lt,  
     movie_companies AS mc,  
     movie_keyword AS mk,  
     movie_link AS ml,  
     title AS t  
WHERE cn.country_code != '[pl]'  
     AND (cn.name LIKE '20th Century Fox%'  
          OR cn.name LIKE 'Twentieth Century Fox%')  
     AND ct.kind != 'production companies'  
     AND ct.kind IS NOT NULL  
     AND k.keyword IN ('sequel',  
                       'revenge',  
                       'based-on-novel')  
  
     AND mc.note IS NOT NULL  
     AND t.production_year > 1950  
     AND lt.id = ml.link_type_id  
     AND ml.movie_id = t.id  
     AND t.id = mk.movie_id  
     AND mk.keyword_id = k.id  
     AND t.id = mc.movie_id  
     AND mc.company_type_id = ct.id  
     AND mc.company_id = cn.id  
     AND ml.movie_id = mk.movie_id  
     AND ml.movie_id = mc.movie_id  
     AND mk.movie_id = mc.movie_id;  
alena=# SET show_node_sign = 'off';  
SET query_inadequate_execution_time = -1;  
SET replan_regression_mode = 'off';  
SET  
SET  
SET  
alena=#
```

```
SET
SET
alena=# EXPLAIN (ANALYZE, COSTS OFF)
SELECT MIN(cn.name) AS from_company,
       MIN(mc.note) AS production_note,
       MIN(t.title) AS movie_based_on_book
FROM   company_name AS cn,
       company_type AS ct,
       keyword AS k,
       link_type AS lt,
       movie_companies AS mc,
       movie_keyword AS mk,
       movie_link AS ml,
       title AS t
WHERE  cn.country_code != '[pl]'
      AND (cn.name LIKE '20th Century Fox%'
          OR cn.name LIKE 'Twentieth Century Fox%')
      AND ct.kind != 'production companies'
      AND ct.kind IS NOT NULL
      AND k.keyword IN ('sequel',
                       'revenge',
                       'based-on-novel')

      AND mc.note IS NOT NULL
      AND t.production_year > 1950
      AND lt.id = ml.link_type_id
      AND ml.movie_id = t.id
      AND t.id = mk.movie_id
      AND mk.keyword_id = k.id
      AND t.id = mc.movie_id
      AND mc.company_type_id = ct.id
      AND mc.company_id = cn.id
      AND ml.movie_id = mk.movie_id
      AND ml.movie_id = mc.movie_id
      AND mk.movie_id = mc.movie_id;
```

## QUERY PLAN

```

-----
Aggregate (actual time=4536.445..4537.786 rows=1 loops=1)
->  Nested Loop (actual time=132.338..4533.999 rows=6946 loops=1)
    ->  Gather (actual time=132.322..4519.439 rows=6946 loops=1)
        Workers Planned: 2
        Workers Launched: 2
        ->  Nested Loop (actual time=95.081..1774.058 rows=2315 loops=3)
            Join Filter: (ml.movie_id = t.id)
            ->  Nested Loop (actual time=95.018..1759.293 rows=2315 loops=3)
                ->  Nested Loop (actual time=5.801..435.827 rows=260259 loops=3)
                    ->  Merge Join (actual time=5.776..249.789 rows=260316 loops=3)
                        Merge Cond: (mk.movie_id = ml.movie_id)
                        ->  Nested Loop (actual time=4.164..50.417 rows=4026 loops=3)
                            ->  Nested Loop (actual time=3.408..46.760 rows=199 loops=3)
                                ->  Parallel Index Scan using movie_id_movie_companies on movie_companies mc (ac
                                tual time=0.047..24.329 rows=17309 loops=3)
                                    Filter: (note IS NOT NULL)
                                    Rows Removed by Filter: 16835
                                ->  Memoize (actual time=0.001..0.001 rows=0 loops=51926)
                                    Cache Key: mc.company_id
                                    Cache Mode: logical
                                    Hits: 14850  Misses: 1759  Evictions: 0  Overflows: 0  Memory Usage: 118kB
                                    Worker 0:  Hits: 1138  Misses: 257  Evictions: 0  Overflows: 0  Memory Usa
                                    ge: 18kB
                                    Worker 1:  Hits: 31080  Misses: 2842  Evictions: 0  Overflows: 0  Memory U
                                    sage: 190kB
                                ->  Index Scan using company_name_pkey on company_name cn (actual time=0.0
                                07..0.007 rows=0 loops=4858)
                                    Index Cond: (id = mc.company_id)
                                    Filter: (((country_code)::text <> '[pl]')::text) AND ((name ~~ '20th
Century Fox%'::text) OR (name ~~ 'Twentieth Century Fox%'::text)))

```

```

tmp_install/bin/psql -d alena
File Edit View Search Terminal Help
Index Cond: (id = mc.company_id)
Filter: (((country_code)::text <> '[pl]'::text) AND ((name ~~ '20th
Century Fox%'::text) OR (name ~~ 'Twentieth Century Fox%'::text)))
Rows Removed by Filter: 1
-> Index Scan using movie_id_movie_keyword on movie_keyword mk (actual time=0.006..0.
013 rows=20 loops=598)
Index Cond: (movie_id = mc.movie_id)
-> Index Scan using movie_id_movie_link on movie_link ml (actual time=0.035..114.489 rows=2
89906 loops=3)
-> Memoize (actual time=0.000..0.000 rows=1 loops=780947)
Cache Key: mc.company_type_id
Cache Mode: logical
Hits: 61925 Misses: 2 Evictions: 0 Overflows: 0 Memory Usage: 1kB
Worker 0: Hits: 693055 Misses: 1 Evictions: 0 Overflows: 0 Memory Usage: 1kB
Worker 1: Hits: 25963 Misses: 1 Evictions: 0 Overflows: 0 Memory Usage: 1kB
-> Index Scan using company_type_pkey on company_type ct (actual time=0.015..0.016 rows=1 l
oops=4)
Index Cond: (id = mc.company_type_id)
Filter: ((kind IS NOT NULL) AND ((kind)::text <> 'production companies'::text))
Rows Removed by Filter: 0
-> Index Scan using keyword_pkey on keyword k (actual time=0.005..0.005 rows=0 loops=780776)
Index Cond: (id = mk.keyword_id)
Filter: (keyword = ANY ('{sequel, revenge, based-on-novel}'::text[]))
Rows Removed by Filter: 1
-> Index Scan using title_pkey on title t (actual time=0.006..0.006 rows=1 loops=6946)
Index Cond: (id = mk.movie_id)
Filter: (production_year > 1950)
-> Index Only Scan using link_type_pkey on link_type lt (actual time=0.001..0.001 rows=1 loops=6946)
Index Cond: (id = ml.link_type_id)
Heap Fetches: 0
Planning Time: 16.870 ms
Execution Time: 4538.295 ms
(51 rows)

```

**(END)**

```
keyword AS k,  
link_type AS lt,  
movie_companies AS mc,  
movie_keyword AS mk,  
movie_link AS ml,  
title AS t  
WHERE cn.country_code != '[pl]'  
AND (cn.name LIKE '20th Century Fox%'  
OR cn.name LIKE 'Twentieth Century Fox%')  
AND ct.kind != 'production companies'  
AND ct.kind IS NOT NULL  
AND k.keyword IN ('sequel',  
                  'revenge',  
                  'based-on-novel')  
  
AND mc.note IS NOT NULL  
AND t.production_year > 1950  
AND lt.id = ml.link_type_id  
AND ml.movie_id = t.id  
AND t.id = mk.movie_id  
AND mk.keyword_id = k.id  
AND t.id = mc.movie_id  
AND mc.company_type_id = ct.id  
AND mc.company_id = cn.id  
AND ml.movie_id = mk.movie_id  
AND ml.movie_id = mc.movie_id  
AND mk.movie_id = mc.movie_id;  
alena=# SET show_node_sign = 'on';  
SET query_inadequate_execution_time = 200;  
SET replan_max_attempts = 3;  
SET show_node_sign = 'on';  
SET  
SET  
SET  
SET  
alena=#
```

## QUERY PLAN

```
-----  
Aggregate (actual time=498.356..509.342 rows=1 loops=1)  
  NodeSign: 17171918697763930013  
  -> Nested Loop (actual time=303.648..504.944 rows=6946 loops=1)  
    NodeSign: 6754593112057626571  
    -> Nested Loop (actual time=303.639..486.391 rows=6946 loops=1)  
      NodeSign: 17282434352120576007  
      Join Filter: (ml.movie_id = t.id)  
      -> Nested Loop (actual time=303.626..447.321 rows=6946 loops=1)  
        NodeSign: 3091200912143658509  
        -> Nested Loop (actual time=303.616..432.706 rows=6946 loops=1)  
          NodeSign: 5538086858734275445  
          -> Nested Loop (actual time=300.788..347.784 rows=15510 loops=1)  
            NodeSign: 6884373172202303082  
            Join Filter: (mc.movie_id = ml.movie_id)  
            -> Merge Join (actual time=300.773..326.612 rows=1460 loops=1)  
              NodeSign: 2192450524941596969  
              Merge Cond: (mk.movie_id = ml.movie_id)  
              -> Gather Merge (actual time=300.566..312.811 rows=481 loops=1)  
                NodeSign: 14703440834904157670  
                Workers Planned: 2  
                Workers Launched: 2  
                -> Nested Loop (actual time=2.806..332.433 rows=822 loops=3)  
                  NodeSign: 14703440834904157670  
                  -> Parallel Index Scan using movie_id_movie_keyword on movie_keyword mk (actual  
time=0.045..68.261 rows=166044 loops=3)  
                    NodeSign: 5052084036067167614  
                    -> Memoize (actual time=0.001..0.001 rows=0 loops=498132)  
                      NodeSign: 2320022770352834085  
                      Cache Key: mk.keyword_id  
                      Cache Mode: logical
```

```
tmp_install/bin/psql -d alena
File Edit View Search Terminal Help

NodeSign: 16361454620442577243
-> Index Scan using movie_id_movie_companies on movie_companies mc (actual time=0.006..0.012 rows
=11 loops=1460)
    NodeSign: 7709717219160362211
    Index Cond: (movie_id = mk.movie_id)
    Filter: (note IS NOT NULL)
    Rows Removed by Filter: 2
-> Index Scan using company_name_pkey on company_name cn (actual time=0.005..0.005 rows=0 loops=15510)
    NodeSign: 2709735511316738750
    Index Cond: (id = mc.company_id)
    Filter: (((country_code)::text <> '[pl]')::text) AND ((name ~~ '20th Century Fox%'::text) OR (name
~~ 'Twentieth Century Fox%'::text))
    Rows Removed by Filter: 1
-> Index Only Scan using link_type_pkey on link_type lt (actual time=0.002..0.002 rows=1 loops=6946)
    NodeSign: 18193422119297758337
    Index Cond: (id = ml.link_type_id)
    Heap Fetches: 0
-> Index Scan using title_pkey on title t (actual time=0.005..0.005 rows=1 loops=6946)
    NodeSign: 3401760843123382840
    Index Cond: (id = mk.movie_id)
    Filter: (production_year > 1950)
-> Index Scan using company_type_pkey on company_type ct (actual time=0.002..0.002 rows=1 loops=6946)
    NodeSign: 6709730083376076919
    Index Cond: (id = mc.company_type_id)
    Filter: ((kind IS NOT NULL) AND ((kind)::text <> 'production companies')::text))

Planning Time: 8.819 ms
Execution Time: 509.736 ms
Replan Active: true
Table Entries: 8
Controlled Statements: 1
Replanning Attempts: 2
Total Execution Time: 1043.378 ms
(68 rows)
```

**(END)**

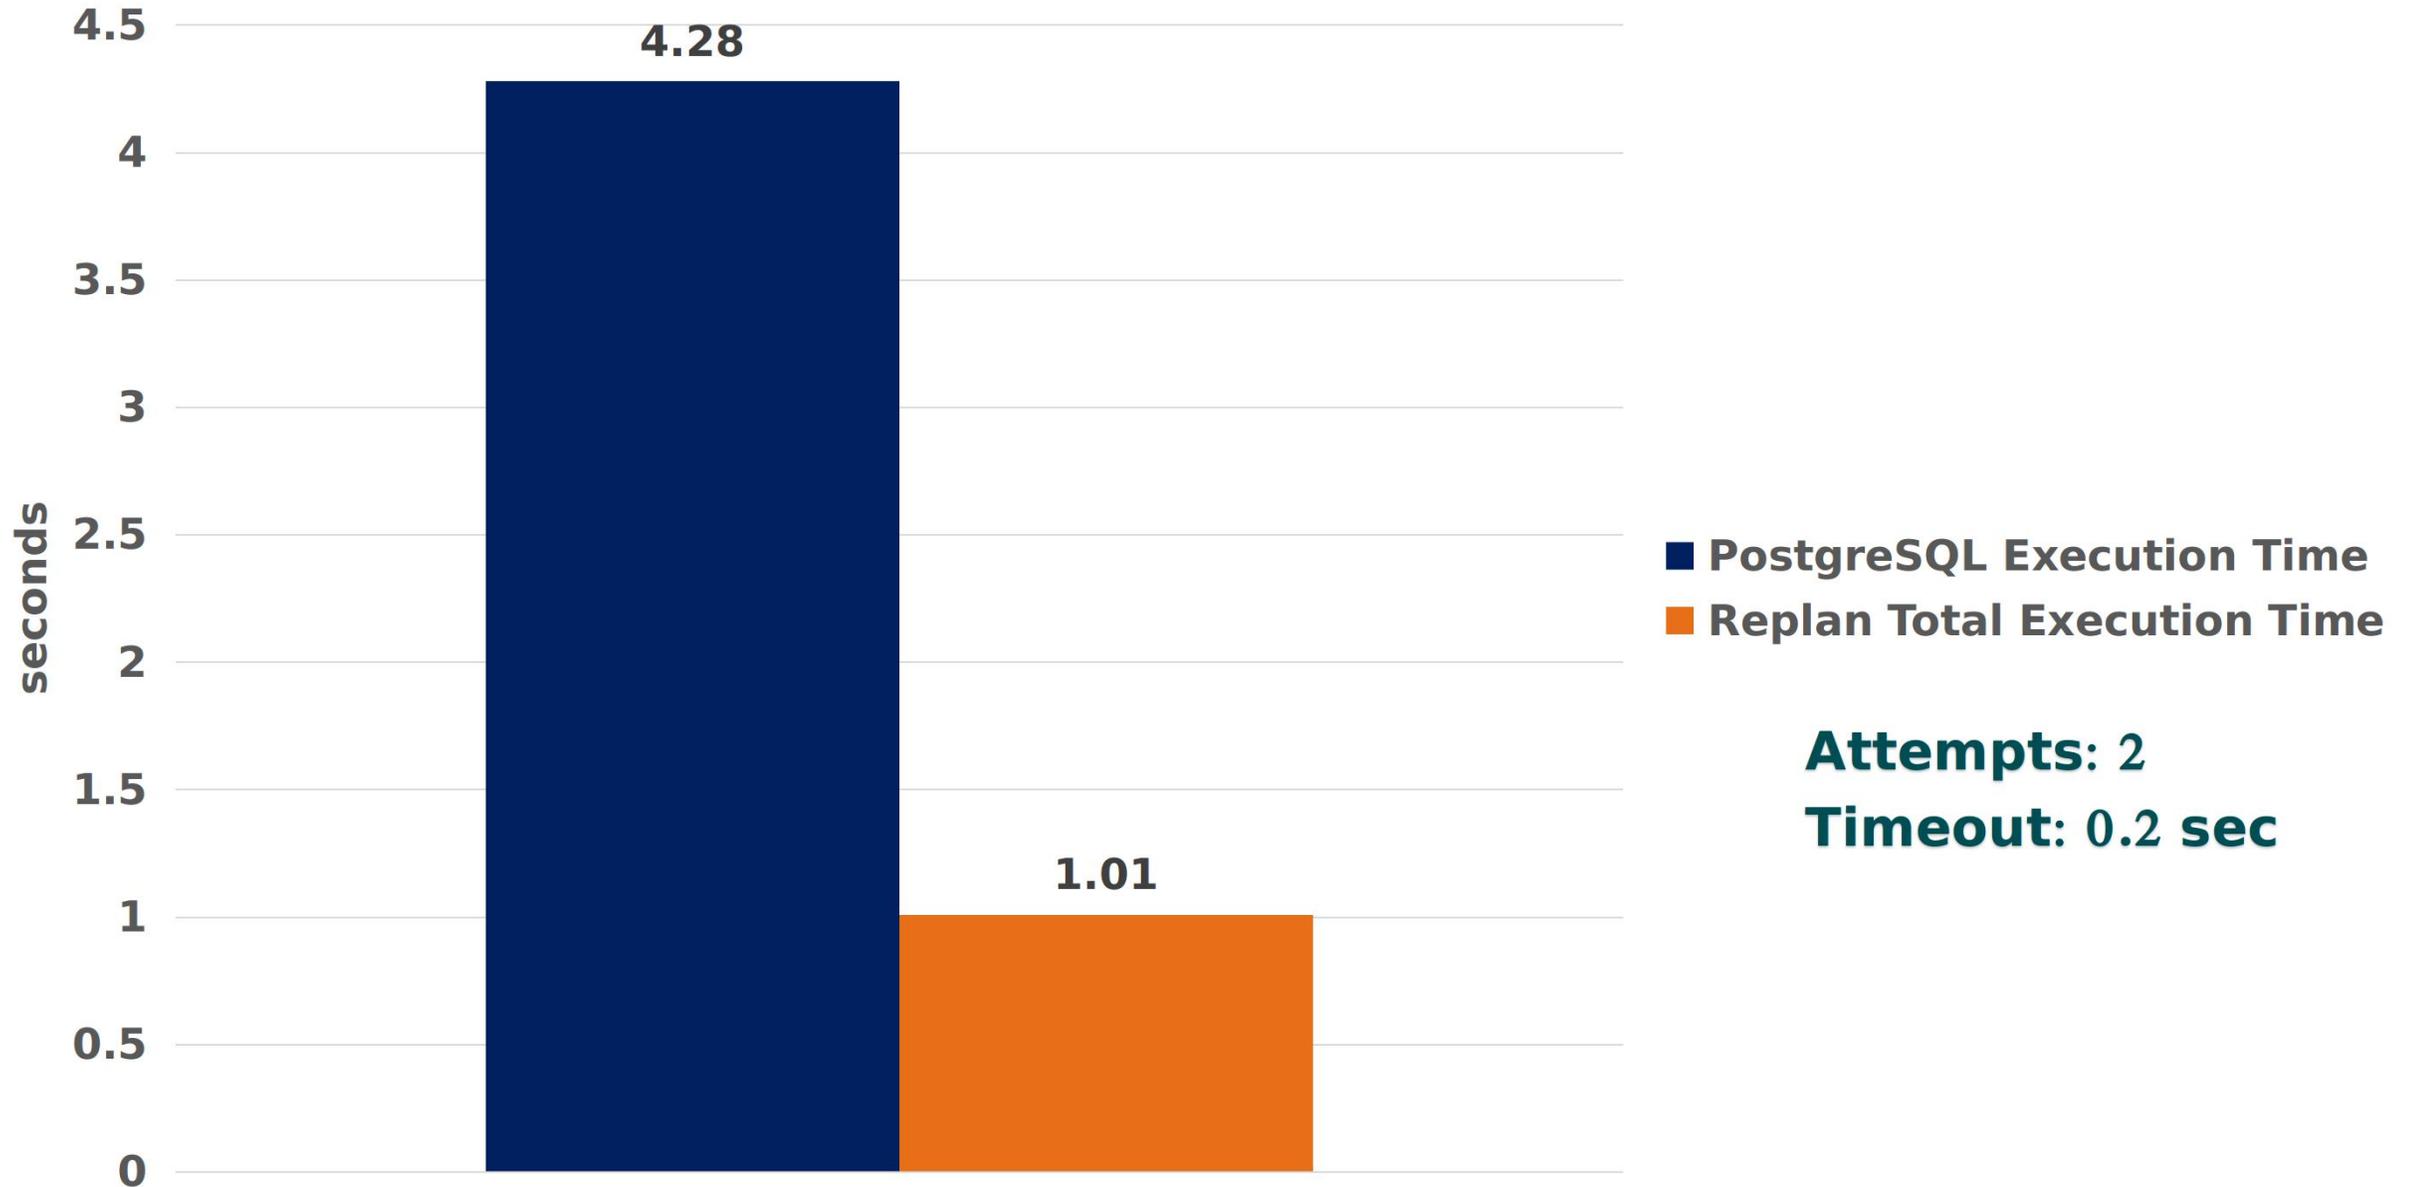
```
tmp_install/bin/psql -d alena
File Edit View Search Terminal Help

NodeSign: 16361454620442577243
-> Index Scan using movie_id_movie_companies on movie_companies mc (actual time=0.006..0.012 rows
=11 loops=1460)
NodeSign: 7709717219160362211
Index Cond: (movie_id = mk.movie_id)
Filter: (note IS NOT NULL)
Rows Removed by Filter: 2
-> Index Scan using company_name_pkey on company_name cn (actual time=0.005..0.005 rows=0 loops=15510)
NodeSign: 2709735511316738750
Index Cond: (id = mc.company_id)
Filter: (((country_code)::text <> '[pl]')::text) AND ((name ~ '20th Century Fox%'::text) OR (name
~~ 'Twentieth Century Fox%'::text))
Rows Removed by Filter: 1
-> Index Only Scan using link_type_pkey on link_type lt (actual time=0.002..0.002 rows=1 loops=6946)
NodeSign: 18193422119297758337
Index Cond: (id = ml.link_type_id)
Heap Fetches: 0
-> Index Scan using title_pkey on title t (actual time=0.005..0.005 rows=1 loops=6946)
NodeSign: 3401760843123382840
Index Cond: (id = mk.movie_id)
Filter: (production_year > 1950)
-> Index Scan using company_type_pkey on company_type ct (actual time=0.002..0.002 rows=1 loops=6946)
NodeSign: 6709730083376076919
Index Cond: (id = mc.company_type_id)
Filter: ((kind IS NOT NULL) AND ((kind)::text <> 'production companies')::text))

Planning Time: 8.819 ms
Execution Time: 509.736 ms
Replan Active: true
Table Entries: 8
Controlled Statements: 1
Replanning Attempts: 2
Total Execution Time: 1043.378 ms
(68 rows)

(END)
```

## Query execution time 11c



```
char_name AS chn,  
cast_info AS ci,  
company_name AS cn,  
info_type AS it,  
info_type AS it3,  
keyword AS k,  
movie_companies AS mc,  
movie_info AS mi,  
movie_keyword AS mk,  
name AS n,  
person_info AS pi,  
role_type AS rt,  
title AS t  
WHERE cct1.kind = 'cast'  
AND cct2.kind = 'complete+verified'  
AND ci.note IN ('(voice)',  
                '(voice: Japanese version)',  
                '(voice) (uncredited)'),  
AND cct2.id = cc.status_id;d_id000 AND 2010  
alena=# SET show_node_sign = 'off';  
SET query_inadequate_execution_time = -1;  
SET replan_regression_mode = 'off';  
SET  
SET  
SET  
alena=#
```



## QUERY PLAN

```
-----  
-----  
Aggregate (actual time=258021.132..258021.140 rows=1 loops=1)  
-> Nested Loop (actual time=386.815..258009.527 rows=5400 loops=1)  
    -> Nested Loop (actual time=1.042..256240.244 rows=189738 loops=1)  
        -> Nested Loop (actual time=0.322..37660.250 rows=41119579 loops=1)  
            -> Nested Loop (actual time=0.298..11614.031 rows=385916 loops=1)  
                -> Nested Loop (actual time=0.282..8378.097 rows=385916 loops=1)  
                    -> Nested Loop (actual time=0.257..1936.360 rows=4028 loops=1)  
                        Join Filter: (kt.id = t.kind_id)  
                        Rows Removed by Join Filter: 8908  
                        -> Seq Scan on kind_type kt (actual time=0.029..0.033 rows=  
1 loops=1)  
                            Filter: ((kind)::text = 'movie'::text)  
                            Rows Removed by Filter: 6  
                        -> Nested Loop (actual time=0.227..1931.584 rows=12936 loop  
s=1)  
                            Join Filter: (cc.movie_id = t.id)  
                            -> Hash Join (actual time=0.104..903.636 rows=61559 l  
oops=1)  
                                Hash Cond: (mi_idx.info_type_id = it2.id)  
                                -> Nested Loop (actual time=0.070..243.104 rows  
=184902 loops=1)
```

:

```
me=0.191..1.567 rows=96 loops=4028)
```

```
Index Cond: (movie_id = t.id)
```

```
-> Index Only Scan using name_pkey on name n (actual time=0.008..0.008
```

```
rows=1 loops=385916)
```

```
Index Cond: (id = ci.person_id)
```

```
Heap Fetches: 0
```

```
-> Index Scan using movie_id_movie_keyword on movie_keyword mk (actual time=0.008..0.047 rows=107 loops=385916)
```

```
Index Cond: (movie_id = t.id)
```

```
-> Index Scan using keyword_pkey on keyword k (actual time=0.005..0.005 rows=0 loops=41119579)
```

```
Index Cond: (id = mk.keyword_id)
```

```
Filter: (keyword = ANY ('{superhero,marvel-comics,based-on-comic,tv-special,fight,violence,magnet,web,claw,laser}'::text[]))
```

```
Rows Removed by Filter: 1
```

```
-> Index Scan using char_name_pkey on char_name chn (actual time=0.009..0.009 rows=0 loops=189738)
```

```
Index Cond: (id = ci.person_role_id)
```

```
Filter: ((name IS NOT NULL) AND ((name ~ '%man%'::text) OR (name ~ '%Man%'::text)))
```

```
)
```

```
Rows Removed by Filter: 0
```

```
Planning Time: 180.249 ms
```

```
Execution Time: 258021.306 ms
```

```
(60 rows)
```

```
(END)
```

```
comp_cast_type AS cct1,  
comp_cast_type AS cct2,  
char_name AS chn,  
cast_info AS ci,  
info_type AS it2,  
keyword AS k,  
kind_type AS kt,  
movie_info_idx AS mi_idx,  
movie_keyword AS mk,  
name AS n,  
title AS t  
WHERE cct1.kind = 'cast'  
AND cct2.kind LIKE '%complete%'  
AND chn.name IS NOT NULL  
AND (chn.name LIKE '%man%'  
OR chn.name LIKE '%Mah%')  
AND it2.info = 'rating'  
AND k.keyword IN ('superhero',  
                  'marvel-comics',  
                  'based-on-comic',  
                  'tv-special',  
                  AND it2.id = mi_idx.info_type_id;  
alena=# SET show_node_sign = 'on';  
SET query_inadequate_execution_time = 1000;  
SET replan_max_attempts = 49;  
SET show_node_sign = 'on';
```

```
        MIN(mi_idx.info) AS rating,
        MIN(t.title) AS complete_hero_movie
FROM complete_cast AS cc,
     comp_cast_type AS cct1,
     comp_cast_type AS cct2,
     char_name AS chn,
     cast_info AS ci,
     info_type AS it2,
     keyword AS k,
     kind_type AS kt,
     movie_info_idx AS mi_idx,
     movie_keyword AS mk,
     name AS n,
     title AS t
WHERE cct1.kind = 'cast'
     AND cct2.kind LIKE '%complete%'
     AND chn.name IS NOT NULL
     AND (chn.name LIKE '%man%'
          OR chn.name LIKE '%Man%')
     AND it2.info = 'rating'
     AND k.keyword IN ('superhero',
                      'marvel-comics',
                      'based-on-comic',
                      'tv-special',
                      'tv-special')
     AND it2.id = mi_idx.info_type_id;
```

## QUERY PLAN

```

-----
Aggregate (actual time=1575.511..1575.520 rows=1 loops=1)
  NodeSign: 15000387547384983206
  -> Hash Join (actual time=132.785..1571.902 rows=5400 loops=1)
    NodeSign: 1832305143051193541
    Hash Cond: (mi_idx.info_type_id = it2.id)
    -> Nested Loop (actual time=132.715..1566.951 rows=16596 loops=1)
      NodeSign: 3668891785607939735
      Join Filter: (mi_idx.movie_id = t.id)
      -> Nested Loop (actual time=132.688..1520.303 rows=5406 loops=1)
        NodeSign: 6566913181348039524
        -> Nested Loop (actual time=132.668..1475.279 rows=5406 loops=1)
          NodeSign: 12481127197580843281
          -> Nested Loop (actual time=130.515..797.260 rows=189844 loops=1)
            NodeSign: 2935371501653627448
            Join Filter: (ci.movie_id = t.id)
            -> Hash Join (actual time=130.485..254.428 rows=1267 loops=1)
              NodeSign: 5564084061457495311
              Hash Cond: (t.kind_id = kt.id)
              -> Nested Loop (actual time=130.469..153.461 rows=1267 loops=1)

```

-> Index Only Scan using name\_pkey on name n (actual time=0.008..0.008 rows=1

loops=5406)

NodeSign: 11336545220415608477

Index Cond: (id = ci.person\_id)

Heap Fetches: 0

-> Index Scan using movie\_id\_movie\_info\_idx on movie\_info\_idx mi\_idx (actual time=0.006..0.007 rows=3 loops=5406)

NodeSign: 4574594900255450271

Index Cond: (movie\_id = ci.movie\_id)

-> Hash (actual time=0.059..0.060 rows=1 loops=1)

NodeSign: 10579708320778064610

Buckets: 1024 Batches: 1 Memory Usage: 9kB

-> Seq Scan on info\_type it2 (actual time=0.052..0.054 rows=1 loops=1)

NodeSign: 10579708320778064610

Filter: ((info)::text = 'rating'::text)

Rows Removed by Filter: 112

Planning Time: 256.081 ms

Execution Time: 1575.765 ms

Replan Active: true

Table Entries: 12

Controlled Statements: 1

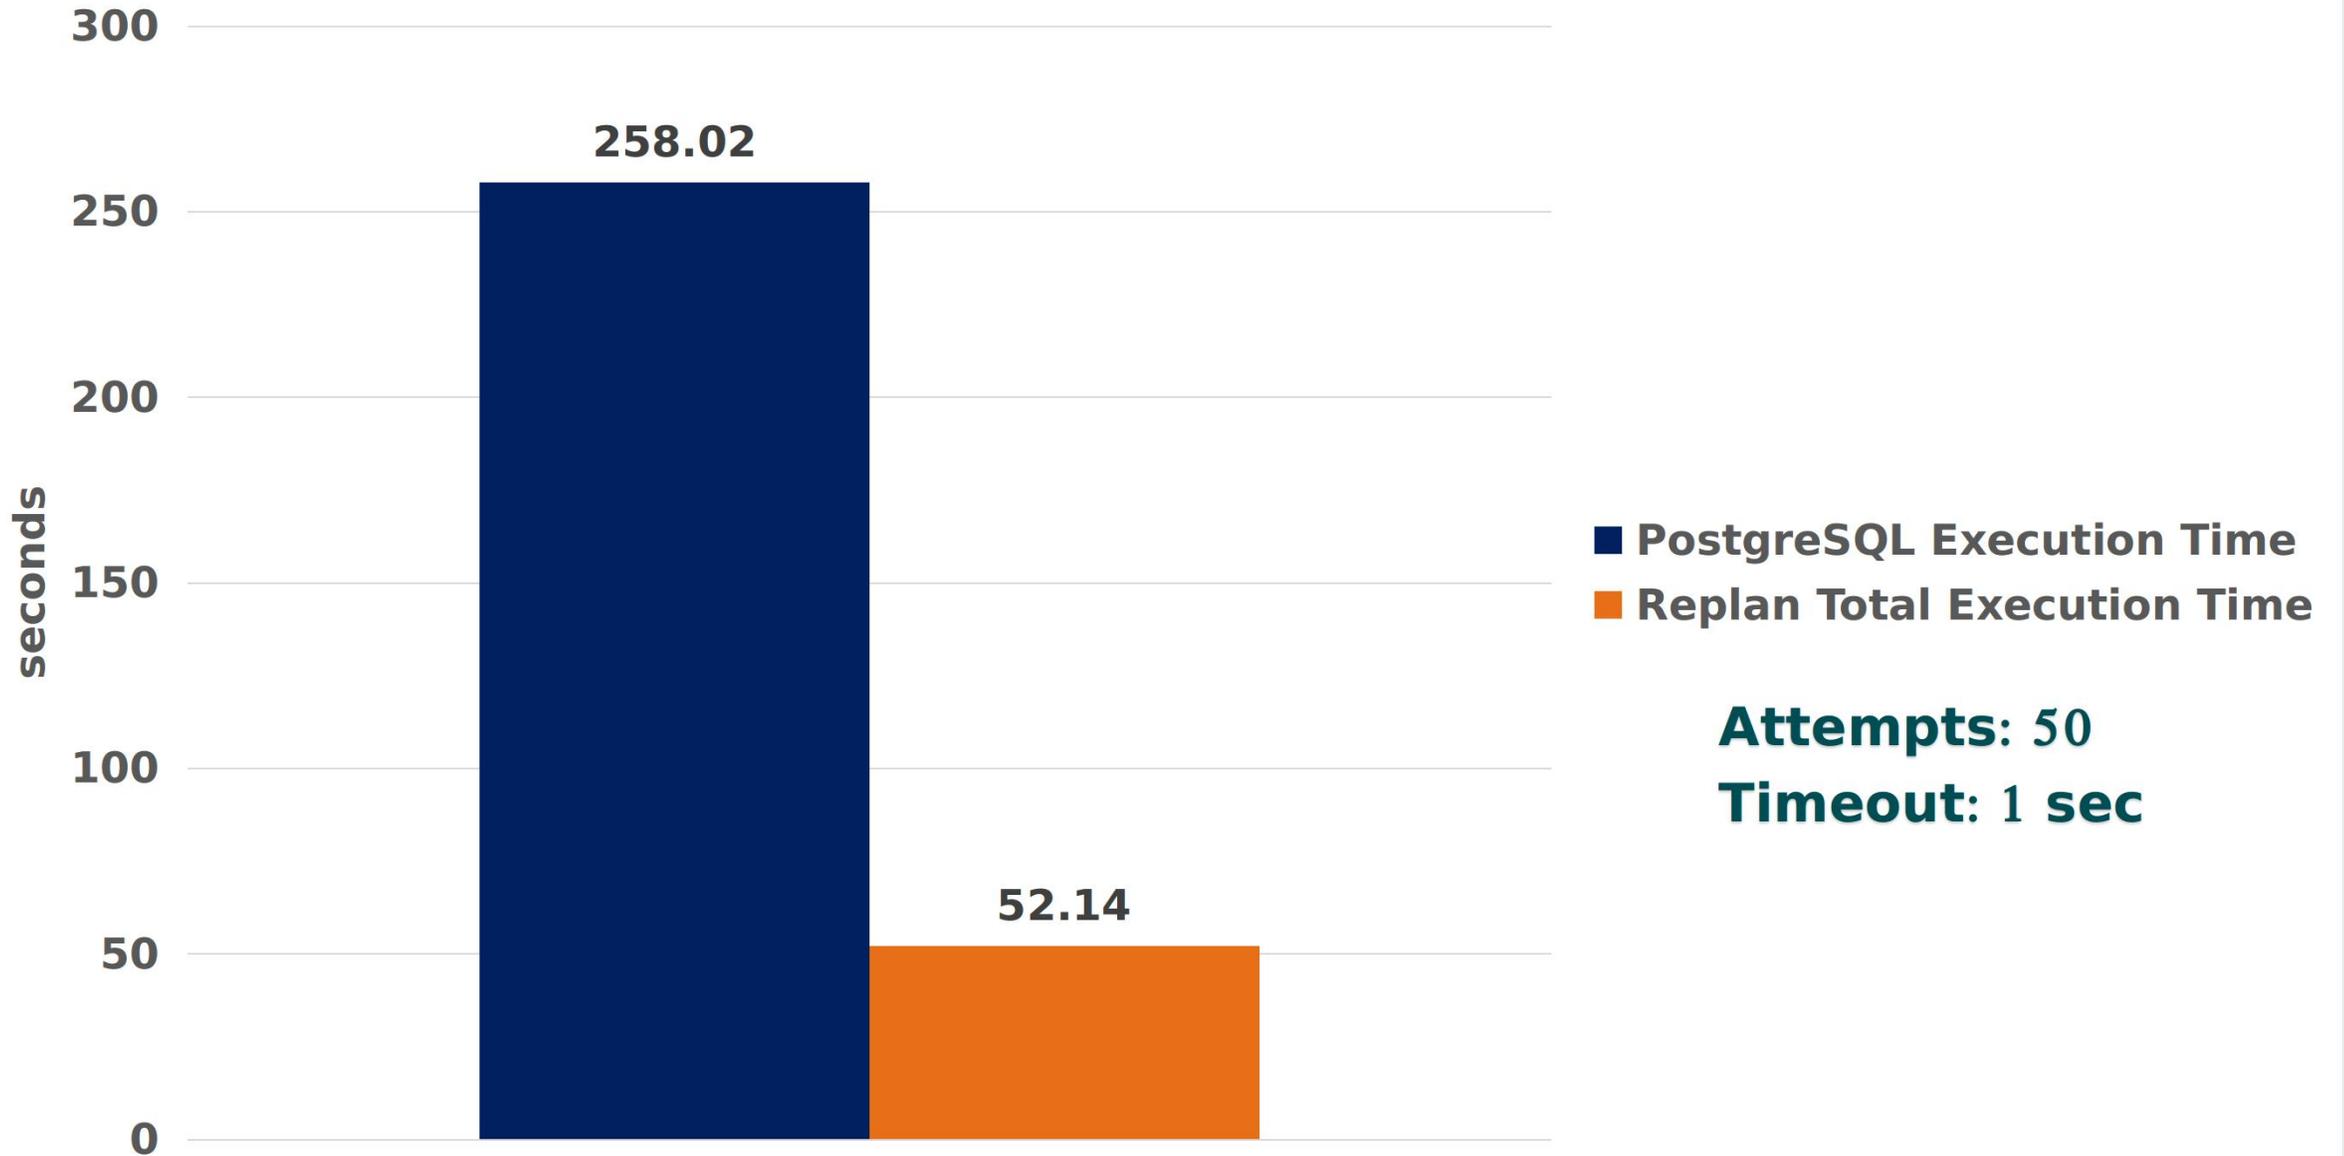
Replanning Attempts: 50

Total Execution Time: 52143.726 ms

(102 rows)

**(END)**

## Query execution time 26c



```
char_name AS chn,  
cast_info AS ci,  
company_name AS cn,  
info_type AS it,  
info_type AS it3,  
keyword AS k,  
movie_companies AS mc,  
movie_info AS mi,  
movie_keyword AS mk,  
name AS n,  
person_info AS pi,  
role_type AS rt,  
title AS t  
WHERE cct1.kind = 'cast'  
AND cct2.kind = 'complete+verified'  
AND ci.note IN ('(voice)',  
                '(voice: Japanese version)',  
                '(voice) (uncredited)',  
                '(voice: English version)')  
AND cn.country_code = '[us]'  
AND it.info = 'release dates'  
AND it3.info = 'trivia'  
AND k.keyword = 'computer-animation'  
AND mi.info IS NOT NULL  
AND (mi.info LIKE 'Japan:%200%'  
     OR mi.info LIKE 'USA:%200%')  
AND n.gender = 'f'  
AND cct2.id = cc.status_id;d_id000 AND 2010  
alena=# SET show_node_sign = 'off';  
SET query_inadequate_execution_time = -1;  
SET replan_regression_mode = 'off';  
SET  
SET  
SET  
alena=#
```

```
        MIN(n.name) AS voicing_actress,
        MIN(t.title) AS voiced_animation
FROM aka_name AS an,
     complete_cast AS cc,
     comp_cast_type AS cct1,
     comp_cast_type AS cct2,
     char_name AS chn,
     cast_info AS ci,
     company_name AS cn,
     info_type AS it,
     info_type AS it3,
     keyword AS k,
     movie_companies AS mc,
     movie_info AS mi,
     movie_keyword AS mk,
     name AS n,
     person_info AS pi,
     role_type AS rt,
     title AS t
WHERE cct1.kind = 'cast'
     AND cct2.kind = 'complete+verified'
     AND ci.note IN ('(voice)',
                    '(voice: Japanese version)',
                    '(voice) (uncredited)',
                    '(voice: English version)')
     AND cn.country_code = '[us]'
     AND it.info = 'release dates'
     AND it3.info = 'trivia'
     AND k.keyword = 'computer-animation'
     AND mi.info IS NOT NULL
     AND (mi.info LIKE 'Japan:%200%'
          OR mi.info LIKE 'USA:%200%')
     AND n.gender = 'f'
     AND cct2.id = cc.status_id;d_id000 AND 2010
```

## QUERY PLAN

```

-----
Aggregate (actual time=3648.462..3648.471 rows=1 loops=1)
-> Nested Loop (actual time=600.959..3639.677 rows=16308 loops=1)
    Join Filter: (cct2.id = cc.status_id)
    Rows Removed by Join Filter: 11714
    -> Nested Loop (actual time=118.619..3588.880 rows=28022 loops=1)
        -> Nested Loop (actual time=118.578..2826.361 rows=149974 loops=1)
            Join Filter: (ci.role_id = rt.id)
            -> Nested Loop (actual time=118.568..2575.913 rows=149974 loops=1)
                Join Filter: (cct1.id = cc.subject_id)
                -> Seq Scan on comp_cast_type cct1 (actual time=0.014..0.016 rows=1 loops=1)
                    Filter: ((kind)::text = 'cast'::text)
                    Rows Removed by Filter: 3
                -> Nested Loop (actual time=118.553..2542.891 rows=149974 loops=1)
                    -> Nested Loop (actual time=81.293..1734.554 rows=169268 loops=1)
                        Join Filter: (mc.movie_id = t.id)
                        -> Nested Loop (actual time=81.279..1587.795 rows=6663 loops=1)
                            Join Filter: (ci.person_id = n.id)
                            -> Nested Loop (actual time=2.232..656.786 rows=175383 loops=1)
                                Join Filter: (it3.id = pi.info_type_id)
                                Rows Removed by Join Filter: 412218
                                -> Seq Scan on info_type it3 (actual time=0.007..0.019 rows=1 loops=1)
                                    Filter: ((info)::text = 'trivia'::text)
                                    Rows Removed by Filter: 112
                                -> Nested Loop (actual time=2.220..579.393 rows=587601 loops=1)
                                    -> Nested Loop (actual time=2.208..182.041 rows=12573 loops=1)
                                        -> Nested Loop (actual time=1.770..141.576 rows=6420 loops=1)
                                            -> Nested Loop (actual time=1.238..94.067 rows=7505 loops=1)
                                                Join Filter: (ci.movie_id = t.id)
                                                -> Nested Loop (actual time=0.375..40.454 rows=357 loop

```

s=1)

:

```

tmp_install/bin/psql -d alena
File Edit View Search Terminal Help
time=0.006..0.006 rows=1 loops=7505)
                                Index Cond: (id = ci.person_role_id)
                                -> Index Only Scan using person_id_aka_name on aka_name an (actual
time=0.005..0.006 rows=2 loops=6420)
                                Index Cond: (person_id = ci.person_id)
                                Heap Fetches: 0
                                -> Index Scan using person_id_person_info on person_info pi (actual time=
0.006..0.022 rows=47 loops=12573)
                                Index Cond: (person_id = ci.person_id)
                                -> Index Scan using name_pkey on name n (actual time=0.005..0.005 rows=0 loops=175383
)
                                Index Cond: (id = pi.person_id)
                                Filter: ((name ~ '%An%'::text) AND ((gender)::text = 'f'::text))
                                Rows Removed by Filter: 1
                                -> Index Scan using movie_id_movie_companies on movie_companies mc (actual time=0.005..0.01
4 rows=25 loops=6663)
                                Index Cond: (movie_id = mk.movie_id)
                                -> Index Scan using movie_id_complete_cast on complete_cast cc (actual time=0.004..0.004 rows=1 l
oops=169268)
                                Index Cond: (movie_id = t.id)
                                -> Seq Scan on role_type rt (actual time=0.001..0.001 rows=1 loops=149974)
                                Filter: ((role)::text = 'actress'::text)
                                Rows Removed by Filter: 1
                                -> Index Scan using company_name_pkey on company_name cn (actual time=0.005..0.005 rows=0 loops=149974)
                                Index Cond: (id = mc.company_id)
                                Filter: ((country_code)::text = '[us]'::text)
                                Rows Removed by Filter: 1
                                -> Seq Scan on comp_cast_type cct2 (actual time=0.001..0.001 rows=1 loops=28022)
                                Filter: ((kind)::text = 'complete+verified'::text)
                                Rows Removed by Filter: 3
Planning Time: 277.900 ms
Execution Time: 3648.654 ms
(84 rows)

```

(END)

```
company_name AS cn,  
info_type AS it,  
info_type AS it3,  
keyword AS k,  
movie_companies AS mc,  
movie_info AS mi,  
movie_keyword AS mk,  
name AS n,  
person_info AS pi,  
role_type AS rt,  
title AS t  
WHERE cct1.kind = 'cast'  
AND cct2.kind = 'complete+verified'  
AND ci.note IN ('(voice)',  
                '(voice: Japanese version)',  
                '(voice) (uncredited)',  
                '(voice: English version)')  
AND cn.country_code = '[us]'  
AND it.info = 'release dates'  
AND it3.info = 'trivia'  
AND k.keyword = 'computer-animation'  
AND mi.info IS NOT NULL  
AND (mi.info LIKE 'Japan:%200%'  
     OR mi.info LIKE 'USA:%200%')  
AND n.gender = 'f'  
AND cct2.id = cc.status_id;d_id000 AND 2010  
alena=# SET show_node_sign = 'on';  
SET query_inadequate_execution_time = 400;  
SET replan_max_attempts = 8;  
SET show_node_sign = 'on';  
SET  
SET  
SET  
SET  
alena=#
```

```
        MIN(n.name) AS voicing_actress,  
        MIN(t.title) AS voiced_animation  
FROM aka_name AS an,  
     complete_cast AS cc,  
     comp_cast_type AS cct1,  
     comp_cast_type AS cct2,  
     char_name AS chn,  
     cast_info AS ci,  
     company_name AS cn,  
     info_type AS it,  
     info_type AS it3,  
     keyword AS k,  
     movie_companies AS mc,  
     movie_info AS mi,  
     movie_keyword AS mk,  
     name AS n,  
     person_info AS pi,  
     role_type AS rt,  
     title AS t  
WHERE cct1.kind = 'cast'  
     AND cct2.kind = 'complete+verified'  
     AND ci.note IN ('(voice)',  
                    '(voice: Japanese version)',  
                    '(voice) (uncredited)',  
                    '(voice: English version)')  
     AND cn.country_code = '[us]'  
     AND it.info = 'release dates'  
     AND it3.info = 'trivia'  
     AND k.keyword = 'computer-animation'  
     AND mi.info IS NOT NULL  
     AND (mi.info LIKE 'Japan:%200%'  
          OR mi.info LIKE 'USA:%200%')  
     AND n.gender = 'f'  
     AND cct2.id = cc.status_id;d_id000 AND 2010
```

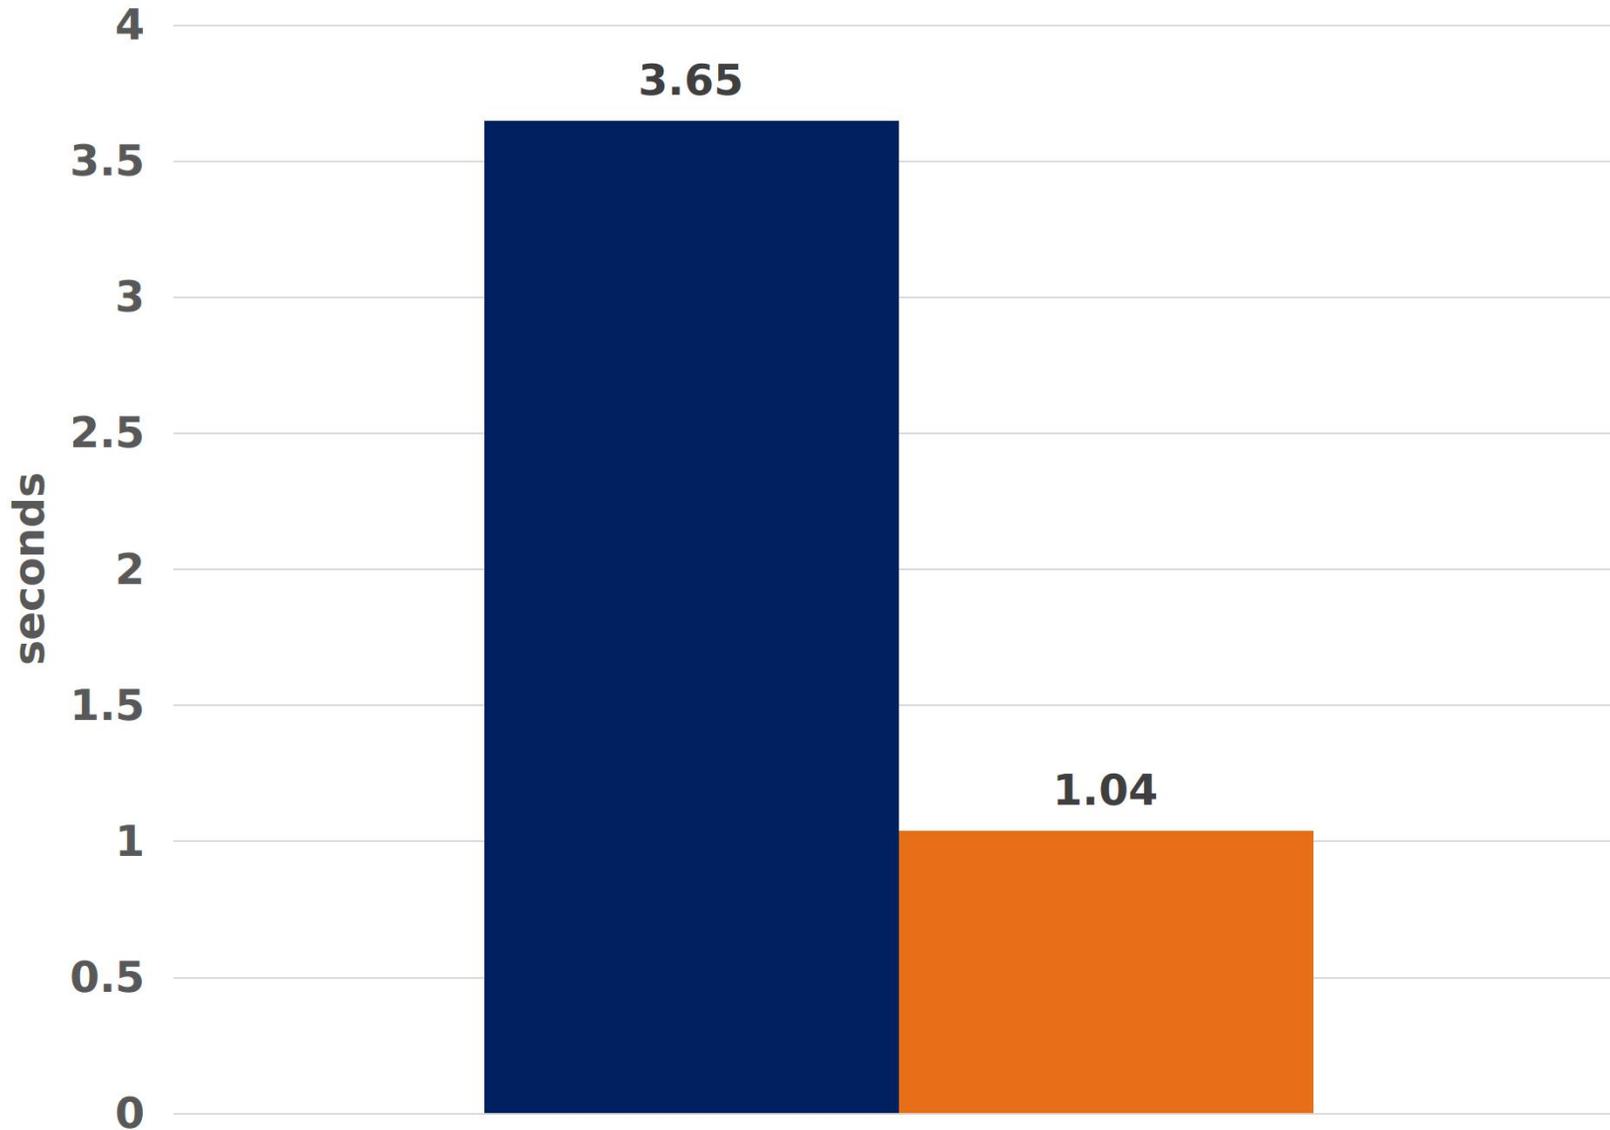
## QUERY PLAN

```
-----  
Aggregate (actual time=115.024..115.035 rows=1 loops=1)  
  NodeSign: 12118314225425016789  
    -> Nested Loop (actual time=13.024..107.882 rows=16308 loops=1)  
      NodeSign: 16222556067297222035  
      Join Filter: (it3.id = pi.info_type_id)  
      Rows Removed by Join Filter: 51519  
      -> Seq Scan on info_type it3 (actual time=0.017..0.031 rows=1 loops=1)  
        NodeSign: 9388138327776621521  
        Filter: ((info)::text = 'trivia'::text)  
        Rows Removed by Filter: 112  
      -> Nested Loop (actual time=12.951..99.561 rows=67827 loops=1)  
        NodeSign: 8132073538093871623  
        Join Filter: (n.id = pi.person_id)  
        -> Nested Loop (actual time=12.935..57.652 rows=189 loops=1)  
          NodeSign: 12680341814512330301  
          Join Filter: (ci.role_id = rt.id)  
          -> Nested Loop (actual time=12.929..57.246 rows=189 loops=1)  
            NodeSign: 9818716129878051414  
            -> Nested Loop (actual time=12.891..52.373 rows=858 loops=1)  
              NodeSign: 14179132961348800920  
              Join Filter: (mc.movie_id = t.id)  
              -> Nested Loop (actual time=12.878..51.508 rows=33 loops=1)  
                NodeSign: 13250385789945565089  
                Join Filter: (an.person_id = n.id)  
                -> Nested Loop (actual time=12.865..51.403 rows=8 loops=1)  
                  NodeSign: 10465605790635078881  
                  Join Filter: (it.id = mi.info_type_id)  
                  -> Nested Loop (actual time=12.855..51.337 rows=8 loops=1)  
                    NodeSign: 17191249967633966537  
                    -> Nested Loop (actual time=6.399..51.255 rows=11 loops=1)
```

```
File Edit View Search Terminal Help
-> Seq Scan on info_type it (actual time=0.006..0.006 rows=1 loops=8)
    NodeSign: 7287576821695634503
    Filter: ((info)::text = 'release dates'::text)
    Rows Removed by Filter: 15
-> Index Only Scan using person_id_aka_name on aka_name an (actual time=0.008..0.010 rows=4
loops=8)
    NodeSign: 2867577293557378872
    Index Cond: (person_id = ci.person_id)
    Heap Fetches: 0
-> Index Scan using movie_id_movie_companies on movie_companies mc (actual time=0.007..0.017 rows
=26 loops=33)
    NodeSign: 9088741564608473149
    Index Cond: (movie_id = ci.movie_id)
-> Index Scan using company_name_pkey on company_name cn (actual time=0.005..0.005 rows=0 loops=858)
    NodeSign: 18322911600766040628
    Index Cond: (id = mc.company_id)
    Filter: ((country_code)::text = '[us]'::text)
    Rows Removed by Filter: 1
-> Seq Scan on role_type rt (actual time=0.001..0.001 rows=1 loops=189)
    NodeSign: 12736067679374244103
    Filter: ((role)::text = 'actress'::text)
    Rows Removed by Filter: 1
-> Index Scan using person_id_person_info on person_info pi (actual time=0.012..0.135 rows=359 loops=189)
    NodeSign: 9569043262522494575
    Index Cond: (person_id = ci.person_id)
Planning Time: 281.065 ms
Execution Time: 115.255 ms
Replan Active: true
Table Entries: 17
Controlled Statements: 1
Replanning Attempts: 2
Total Execution Time: 1041.193 ms
(125 rows)
```

(END)

# Query execution time: 29c



■ PostgreSQL Execution Time  
■ Replan Total Execution Time

**Attempts: 2**

**Timeout: 0.4 sec**

```
FROM company_name AS cn,  
     company_type AS ct,  
     keyword AS k,  
     link_type AS lt,  
     movie_companies AS mc,  
     movie_keyword AS mk,  
     movie_link AS ml,  
     title AS t  
WHERE cn.country_code != '[pl]'  
     AND (cn.name LIKE '20th Century Fox%'  
         OR cn.name LIKE 'Twentieth Century Fox%')  
     AND ct.kind != 'production companies'  
     AND ct.kind IS NOT NULL  
     AND k.keyword IN ('sequel',  
                       'revenge',  
                       'based-on-novel')  
  
     AND mc.note IS NOT NULL  
     AND t.production_year > 1950  
     AND lt.id = ml.link_type_id  
     AND ml.movie_id = t.id  
     AND t.id = mk.movie_id  
     AND mk.keyword_id = k.id  
     AND t.id = mc.movie_id  
     AND mc.company_type_id = ct.id  
     AND mc.company_id = cn.id  
     AND ml.movie_id = mk.movie_id  
     AND ml.movie_id = mc.movie_id  
     AND mk.movie_id = mc.movie_id;  
alena=# SET show_node_sign = 'off';  
SET query_inadequate_execution_time = -1;  
SET replan_regression_mode = 'off';  
SET  
SET  
SET  
alena=#
```

```
        MIN(mi_idx.info) AS movie_votes,
        MIN(n.name) AS writer,
        MIN(t.title) AS violent_liongate_movie
FROM cast_info AS ci,
     company_name AS cn,
     info_type AS it1,
     info_type AS it2,
     keyword AS k,
     movie_companies AS mc,
     movie_info AS mi,
     movie_info_idx AS mi_idx,
     movie_keyword AS mk,
     name AS n,
     title AS t
WHERE ci.note IN ('(writer)',
                 '(head writer)',
                 '(written by)',
                 '(story)',
                 '(story editor)')
AND cn.name LIKE 'Lionsgate%'
AND it1.info = 'genres'
AND it2.info = 'votes'
AND k.keyword IN ('murder',
                 'violence',
                 'blood',
                 'gore',
                 'death',
                 'female-nudity',
                 'hospital')
AND mi.info IN ('Horror',
               'Action',
               'Sci-Fi',
               'Thriller',
               'Drama',
               'Comedy',
               'Mystery',
               'Romance',
               'War',
               'History',
               'Biography',
               'Documentary',
               'Sport',
               'Fantasy',
               'Adventure',
               'Animation',
               'Crime',
               'Family',
               'Music',
               'Short',
               'TV-14',
               'TV-MA',
               'TV-PG',
               'TV-Y7-FV',
               'TV-Y')
AND cn.id = mc.company_id;ype_idd
```

## QUERY PLAN

```

-----
Aggregate (actual time=23524.039..23524.149 rows=1 loops=1)
-> Nested Loop (actual time=121.892..23518.337 rows=2825 loops=1)
    Join Filter: (mi.movie_id = t.id)
    -> Nested Loop (actual time=121.854..23485.026 rows=2825 loops=1)
        -> Nested Loop (actual time=121.803..23446.047 rows=2825 loops=1)
            -> Nested Loop (actual time=16.529..17873.626 rows=589677 loops=1)
                Join Filter: (mc.movie_id = mi.movie_id)
                -> Nested Loop (actual time=16.502..16610.476 rows=42900 loops=1)
                    Join Filter: (ci.movie_id = mi.movie_id)
                    -> Gather (actual time=14.493..1527.944 rows=63386 loops=1)
                        Workers Planned: 2
                        Workers Launched: 2
                        -> Nested Loop (actual time=10.951..2618.553 rows=21129 loops=3)
                            -> Nested Loop (actual time=10.920..2499.655 rows=22577 loops=3)
                                Join Filter: (mi.movie_id = mi_idx.movie_id)
                                -> Hash Join (actual time=10.853..650.270 rows=21234 loops=3)
                                    Hash Cond: (mi_idx.info_type_id = it2.id)
                                    -> Nested Loop (actual time=10.584..617.730 rows=63896 loops=3)
                                        -> Nested Loop (actual time=10.515..245.930 rows=25571 loops=3)
                                            -> Parallel Index Scan using keyword_pkey on keyword k (actual
time=0.488..84.619 rows=2 loops=3)
                                                Filter: (keyword = ANY ('{murder,violence,blood,gore,dea
th,female-nudity,hospital}'::text[]))
                                                Rows Removed by Filter: 44721
                                            -> Bitmap Heap Scan on movie_keyword mk (actual time=7.223..6
5.153 rows=10959 loops=7)
                                                Recheck Cond: (k.id = keyword_id)
                                                Heap Blocks: exact=12389
                                        -> Bitmap Index Scan on keyword_id_movie_keyword (actual
time=3.681..3.681 rows=10959 loops=7)

```

```

File Edit View Search Terminal Help
=3)
                                Filter: ((info)::text = 'votes'::text)
                                Rows Removed by Filter: 112
-> Index Scan using movie_id_movie_info on movie_info mi (actual time=0.055..0.
086 rows=1 loops=63701)
                                Index Cond: (movie_id = mk.movie_id)
                                Filter: (info = ANY ('{Horror,Action,Sci-Fi,Thriller,Crime,War}'::text[]))
                                Rows Removed by Filter: 58
-> Index Scan using info_type_pkey on info_type it1 (actual time=0.004..0.004 rows=1
loops=67732)
                                Index Cond: (id = mi.info_type_id)
                                Filter: ((info)::text = 'genres'::text)
                                Rows Removed by Filter: 0
-> Index Scan using movie_id_cast_info on cast_info ci (actual time=0.187..0.237 rows=1 loops=633
86)
                                Index Cond: (movie_id = mk.movie_id)
                                Filter: (note = ANY ('{(writer),"(head writer)","(written by)",(story),"(story editor)"}'::t
ext[]))
                                Rows Removed by Filter: 63
-> Index Scan using movie_id_movie_companies on movie_companies mc (actual time=0.012..0.022 rows=14 lo
ops=42900)
                                Index Cond: (movie_id = mk.movie_id)
-> Index Scan using company_name_pkey on company_name cn (actual time=0.009..0.009 rows=0 loops=589677)
                                Index Cond: (id = mc.company_id)
                                Filter: (name ~ 'Lionsgate%'::text)
                                Rows Removed by Filter: 1
-> Index Scan using name_pkey on name n (actual time=0.012..0.012 rows=1 loops=2825)
                                Index Cond: (id = ci.person_id)
-> Index Scan using title_pkey on title t (actual time=0.010..0.010 rows=1 loops=2825)
                                Index Cond: (id = mk.movie_id)
Planning Time: 111.464 ms
Execution Time: 23524.423 ms
(58 rows)

```

(END)

```
FROM cast_info AS ci,  
     company_name AS cn,  
     info_type AS it1,  
     info_type AS it2,  
     keyword AS k,  
     movie_companies AS mc,  
     movie_info AS mi,  
     movie_info_idx AS mi_idx,  
     movie_keyword AS mk,  
     name AS n,  
     title AS t  
WHERE ci.note IN ('(writer)',  
                 '(head writer)',  
                 '(written by)',  
                 '(story)',  
                 '(story editor)')  
  
AND cn.name LIKE 'Lionsgate%'  
AND it1.info = 'genres'  
AND it2.info = 'votes'  
AND k.keyword IN ('murder',  
                 'violence',  
                 'blood',  
                 'gore',  
                 'death',  
                 'female-nudity',  
                 'hospital')  
  
AND mi.info IN ('Horror',  
               'Action',  
               'Sci-Fi',  
               'Thriller',  
               'Drama')  
  
AND cn.id = mc.company_id;type idd  
alena=# SET show_node_sign = 'on';  
SET query_inadequate_execution_time = 1000;  
SET replan_max_attempts = 49;  
SET show_node_sign = 'on';
```

```
        MIN(mi_idx.info) AS movie_votes,
        MIN(n.name) AS writer,
        MIN(t.title) AS violent_liongate_movie
FROM cast_info AS ci,
     company_name AS cn,
     info_type AS it1,
     info_type AS it2,
     keyword AS k,
     movie_companies AS mc,
     movie_info AS mi,
     movie_info_idx AS mi_idx,
     movie_keyword AS mk,
     name AS n,
     title AS t
WHERE ci.note IN ('(writer)',
                 '(head writer)',
                 '(written by)',
                 '(story)',
                 '(story editor)')

AND cn.name LIKE 'Lionsgate%'
AND it1.info = 'genres'
AND it2.info = 'votes'
AND k.keyword IN ('murder',
                 'violence',
                 'blood',
                 'gore',
                 'death',
                 'female-nudity',
                 'hospital')

AND mi.info IN ('Horror',
               'Action',
               'Sci-Fi',
               'Thriller',

AND cn.id = mc.company_id;ype_idd
```

## QUERY PLAN

```
-----  
Finalize Aggregate (actual time=993.559..997.579 rows=1 loops=1)  
  NodeSign: 11218188796493155275  
  -> Gather (actual time=842.415..997.549 rows=2 loops=1)  
    NodeSign: not supported  
    Workers Planned: 1  
    Workers Launched: 1  
    -> Partial Aggregate (actual time=914.130..914.141 rows=1 loops=2)  
      NodeSign: not supported  
      -> Nested Loop (actual time=3.760..912.283 rows=1412 loops=2)  
        NodeSign: 8816215010236051633  
        -> Nested Loop (actual time=3.507..405.838 rows=53703 loops=2)  
          NodeSign: 5123304147640578892  
          -> Nested Loop (actual time=3.441..345.996 rows=508 loops=2)  
            NodeSign: 5343039674426417939  
            Join Filter: (mi.movie_id = t.id)  
            -> Nested Loop (actual time=3.385..339.300 rows=508 loops=2)  
              NodeSign: 8848267569147087229  
              -> Nested Loop (actual time=3.331..330.897 rows=508 loops=2)  
                NodeSign: 550196978246219782  
                Join Filter: (ci.movie_id = mi.movie_id)  
                -> Nested Loop (actual time=2.646..130.598 rows=618 loops=2)  
                  NodeSign: 11883474906559214192  
                  -> Nested Loop (actual time=2.617..126.712 rows=648 loops=2)  
                    NodeSign: 4154385644569374001  
                    Join Filter: (mi.movie_id = mi_idx.movie_id)  
                    -> Hash Join (actual time=1.697..56.300 rows=724 loops=2)  
                      NodeSign: 7922155712431859854  
                      Hash Cond: (mi_idx.info_type_id = it2.id)  
                      -> Nested Loop (actual time=1.428..54.807 rows=2185 loops=2)  
                        NodeSign: 18054089972402277607
```

```
File Edit View Search Terminal Help
Index Cond: (id = mi.info_type_id)
Filter: ((info)::text = 'genres'::text)
Rows Removed by Filter: 0
-> Index Scan using movie_id_cast_info on cast_info ci (actual time=0.239..0.323 rows
=1 loops=1235)
NodeSign: 14131283812421365639
Index Cond: (movie_id = mc.movie_id)
Filter: (note = ANY ('{(writer),"(head writer)","(written by)","(story e
ditor)"}'::text[]))
Rows Removed by Filter: 102
-> Index Scan using name_pkey on name n (actual time=0.015..0.015 rows=1 loops=1017)
NodeSign: 2703763068334705397
Index Cond: (id = ci.person_id)
-> Index Scan using title_pkey on title t (actual time=0.011..0.011 rows=1 loops=1017)
NodeSign: 678843994618044636
Index Cond: (id = mc.movie_id)
-> Index Scan using movie_id_movie_keyword on movie_keyword mk (actual time=0.015..0.080 rows=106 loops
=1017)
NodeSign: 10752057836713899998
Index Cond: (movie_id = t.id)
-> Index Scan using keyword_pkey on keyword k (actual time=0.009..0.009 rows=0 loops=107406)
NodeSign: 13217407192521347227
Index Cond: (id = mk.keyword_id)
Filter: (keyword = ANY ('{murder,violence,blood,gore,death,female-nudity,hospital}'::text[]))
Rows Removed by Filter: 1
Planning Time: 146.741 ms
Execution Time: 997.913 ms
Replan Active: true
Table Entries: 11
Controlled Statements: 1
Replanning Attempts: 1
Total Execution Time: 2010.402 ms
(85 rows)
```

(END)

## Query execution time (31c)

