

# Indexes

# Right Place - Right Time

By Karina Ruzinov and Jonathan Salama

March 20 2019

**TRANSFIX** 



Transfix

# Next Generation Online Freight Brokerage

## Facts

Founded in  
2013

HQ in New York  
with 150  
employees

Backed by NEA,  
Canvas Ventures  
& Lerer-Hippeau  
\$78.5M raised

Named to Forbes  
“Next Billion-Dollar  
Startup List 2018”

**Transfix is a  
marketplace for full  
truckload shipping.**

Acting as a truckload broker, we execute truckload shipments and enable large retailers, CPG companies, and manufacturers to ship with a highly fragmented supply base of independent truck drivers and small trucking companies.

We create value by leveraging mobile technology, data, and AI to automate processes, improve matching, and increase utilization of trucks.

**TRANSFIX**

# Key Customers



**BARNES & NOBLE**



**Casper**

**STAPLES**



*J. Crew*



**L'ORÉAL**



# Hello.



**Jonathan Salama**

CTO & Co-founder, Transfix  
Previously Gilt Group, Microsoft



**Karina Ruzinov**

Software Engineering Manager, Transfix  
Previously Warby Parker, Refinery29

**“A goal without a plan is just a wish”**

*Antoine de St-Exupery*




# Sequential Scan

**Your database can be represented  
by a giant set of drawers**

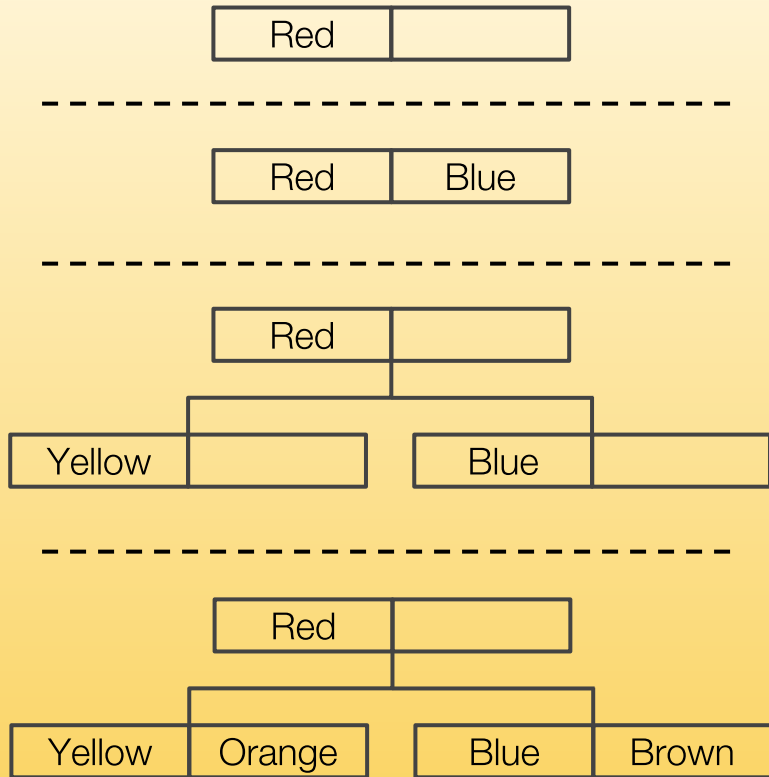
**It's easy to find a pair of socks  
if there are only  
two items in your drawers**

**Things become slightly harder  
if you have millions of items**



# Index Scan

An index could be represented  
as the map to help you  
find your *\*snazzy\** red socks.



# Most Common Index: B TREE

**B is not for Binary**

**Self balancing tree data structure**

**Many other types of indexes:  
Hash, GiST, SP-GiST, GIN and  
BRIN.**

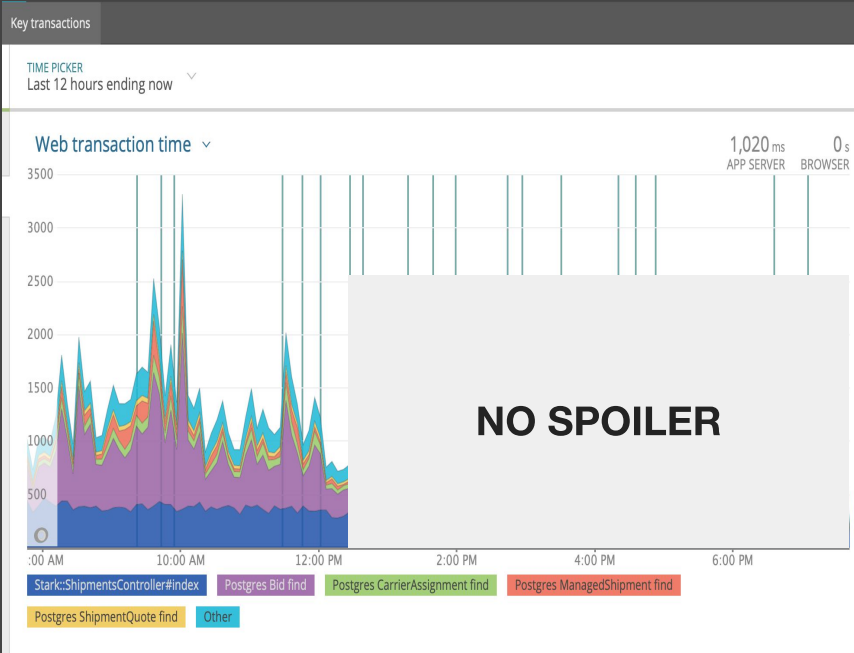


# Our Problem

**Tools we used to identify the issue**

**PSQL is clearly the pain point**

**Where do we go from here**



# Explain & Analyze with PSQL

```
transfix=# explain analyze select * from shipments s join shipment_stops stops on stops.shipment_id = s.id where s.shipper_miles > 10;
```

## QUERY PLAN

```
Merge Join (cost=11.66..226277.87 rows=361110 width=5417) (actual time=0.018..3186.715 rows=369510 loops=1)
  Merge Cond: (s.id = stops.shipment_id)
    -> Seq Scan on shipments s (cost=0.42..79181.20 rows=117477 width=3970) (actual time=0.008..312.337 rows=117744 loops=1)
      Filter: (shipper_miles > '10'::numeric)
      Rows Removed by Filter: 8112
    -> Index Scan using shipment_stops_shipment_id_idx on shipment_stops stops (cost=0.42..142229.09 rows=386888 width=1447) (actual time=0.004..2440.596 rows=386888 loops=1)
Planning time: 43.832 ms
Execution time: 14424.479 ms
(8 rows)
```

# An easier way to read this:

<https://explain.depesz.com/>

#	exclusive	inclusive	rows x	rows	loops	node
1.	0.000	14,752.933	↓ 1.0	3,695,010	1	→ <u>Merge Join</u> (cost=11.66..226,277.87 rows=3,611,100 width=5,417) (actual time=0.018..14,752.933 rows=3,695,010 loops=1) Merge Cond: (s.id = stops.shipment_id)
2.	14,312.337	14,312.337	↓ 10.0	1,177,404	1	→ <u>Seq Scan</u> on shipments s (cost=0.42..791,801.20 rows=117,477 width=3,970) (actual time=0.008..14,312.337 rows=1,177,404 loops=1) Filter: (shipper_miles > '10'::numeric) Rows Removed by Filter: 8112
3.	440.596	440.596	↑ 1.0	386,888	1	→ <u>Index Scan</u> using shipment_stops_shipment_id_idx on shipment_stops stops (cost=0.42..142,229.09 rows=386,888 width=1,447) (actual time=0.004..440.596 rows=386,888 loops=1)

# Remembering the red socks!

Now that we understand our plan we need to create our index on the right column

```
transfix_dev=# create index shipments_shipper_miles on shipments(shipper_miles);  
CREATE INDEX
```



<https://www.postgresql.org/docs/10/sql-createindex.html>

**“Everyone has a plan until they get  
punched in the mouth”**

***Myke Tyson***



# The Pitfalls


**Applying an index can be dangerous**

**Queries don't always follow your plan**

**Unnecessary indexes can waste memory and slow your writes down!**

# Results

Queries goes from 2 to 3 sec to  
few milliseconds!









# Proactively monitoring for performance issues

**Educating your team**

**Empowering your team**

**Continuously monitoring**

# Q & A

A white semi-truck is driving away on a winding road during a sunset. The sun is low on the horizon, casting a warm orange glow over the scene. The sky is filled with soft, dark clouds. The road curves to the right, and a guardrail is visible on the left side. The truck is in the right lane, moving away from the viewer.

**Thank You!**

**TRANSFIX** 