

FULL THROTTLE DATABASE PostgreSQL

WWW.POSTGRESQL.ORG

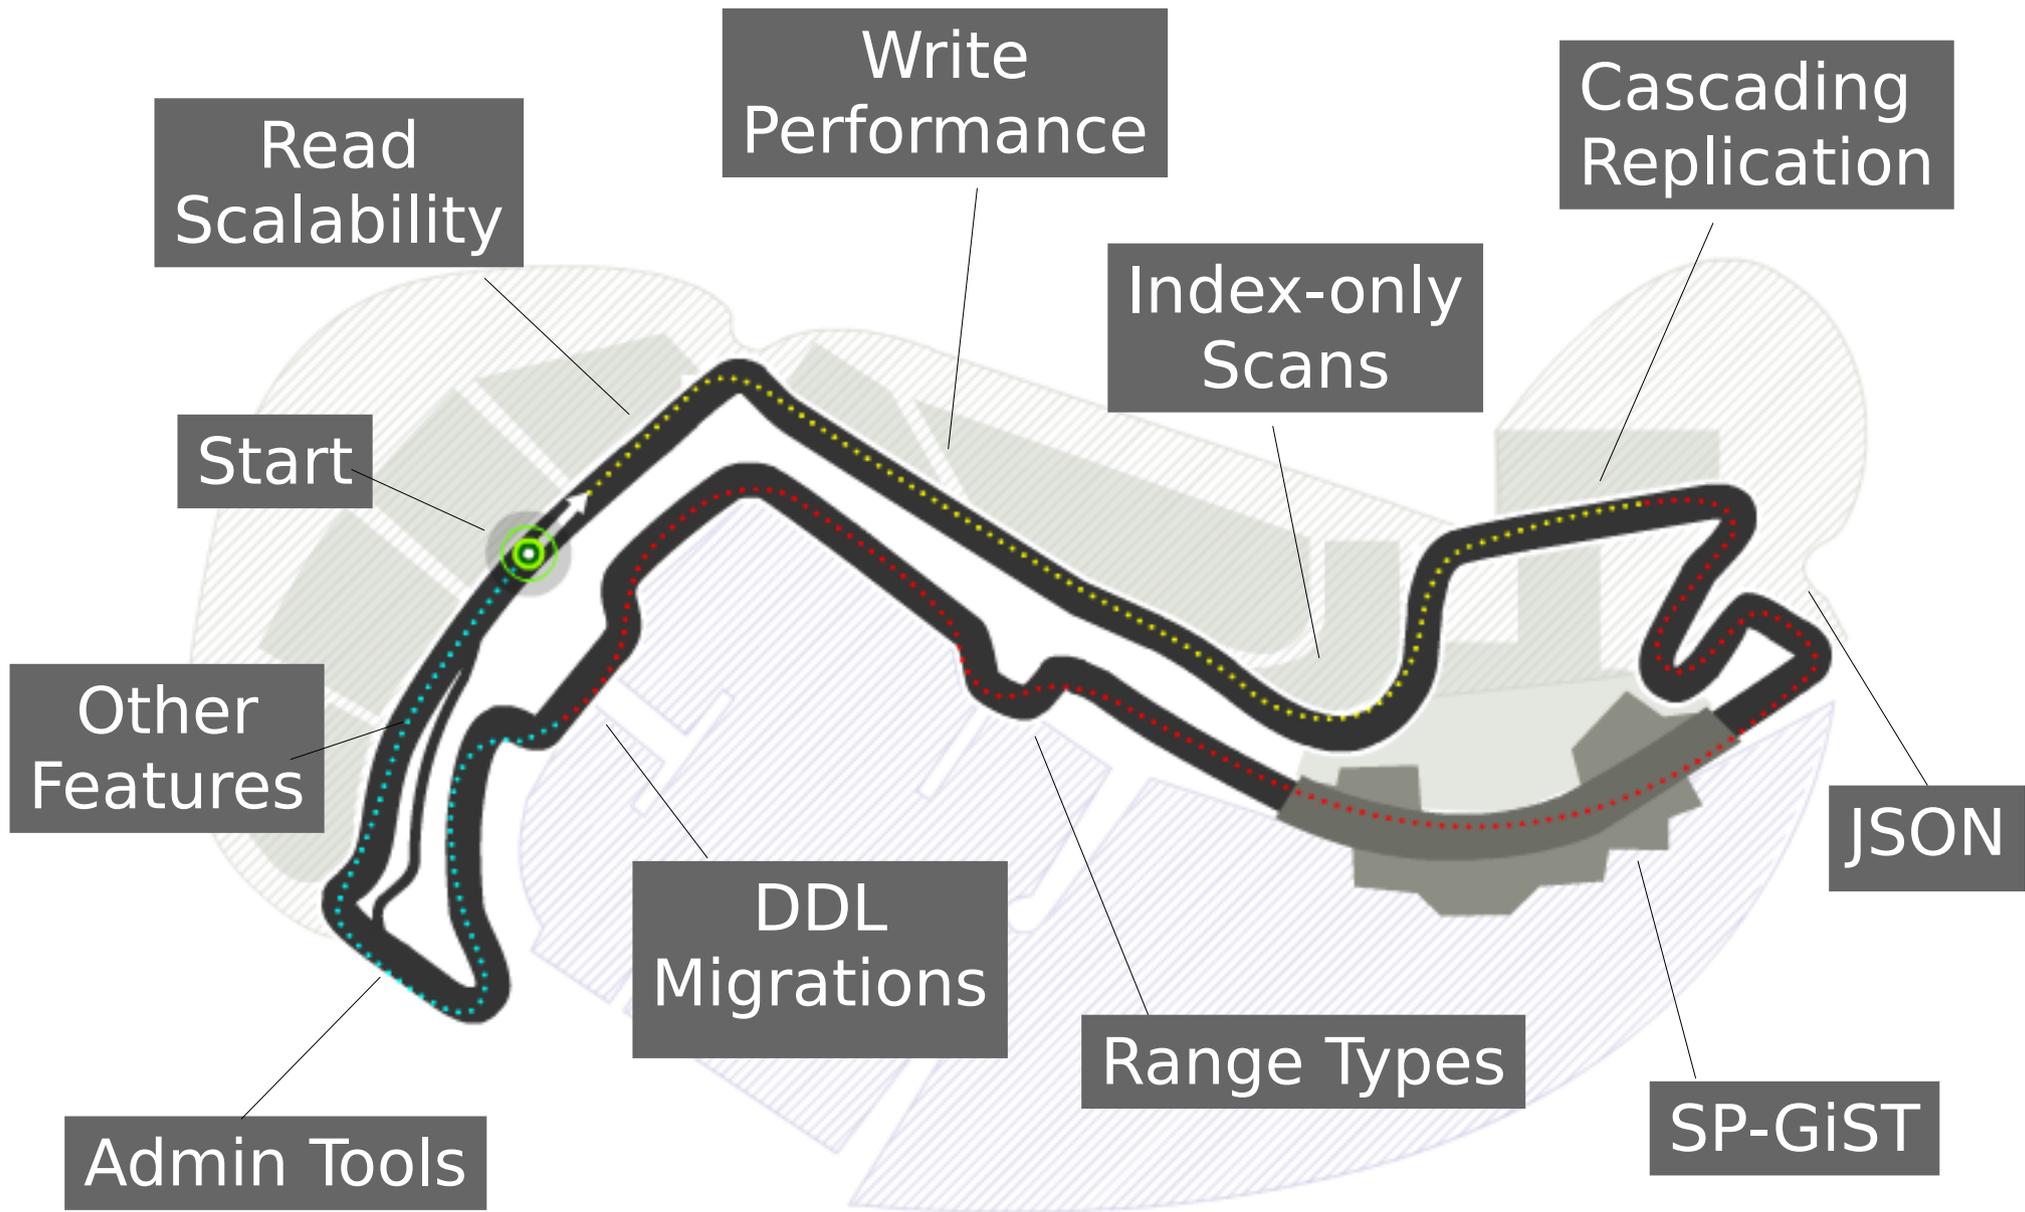
9.2



PGX
POSTGRESQL
EXPERTS, INC.

Chris Lawrence 2012

Josh Berkus
pgCon 2012



New Feature Grand Prix

START YOUR ENGINES!

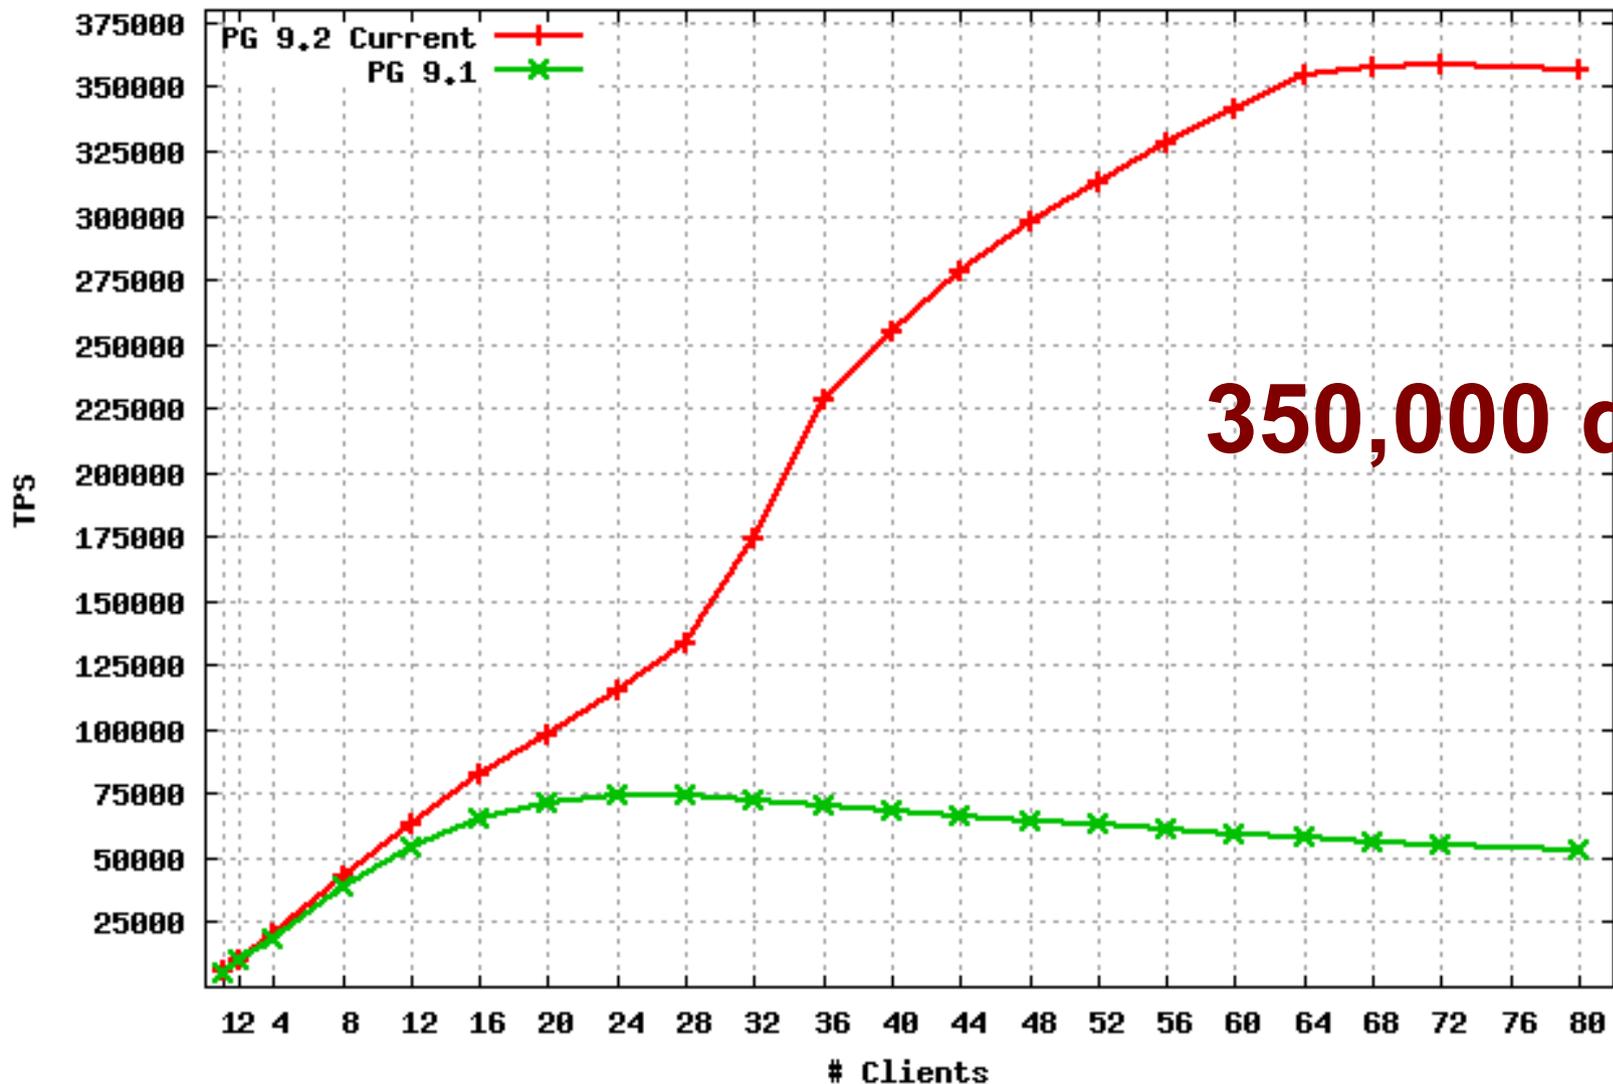


Read Scalability

- fast path sorting
- reduced overhead frequent table locks
- prproc shared memory optimization
- laxy vxid locks
- heap hot search buffer

Read Scalability

pgbench -S, PG 9.2devel as of commit d5881c03
8 x 8-core AMD 6272 Processors
median of 3 5-minute runs, max_connections = 100, shared_buffers = 8GB



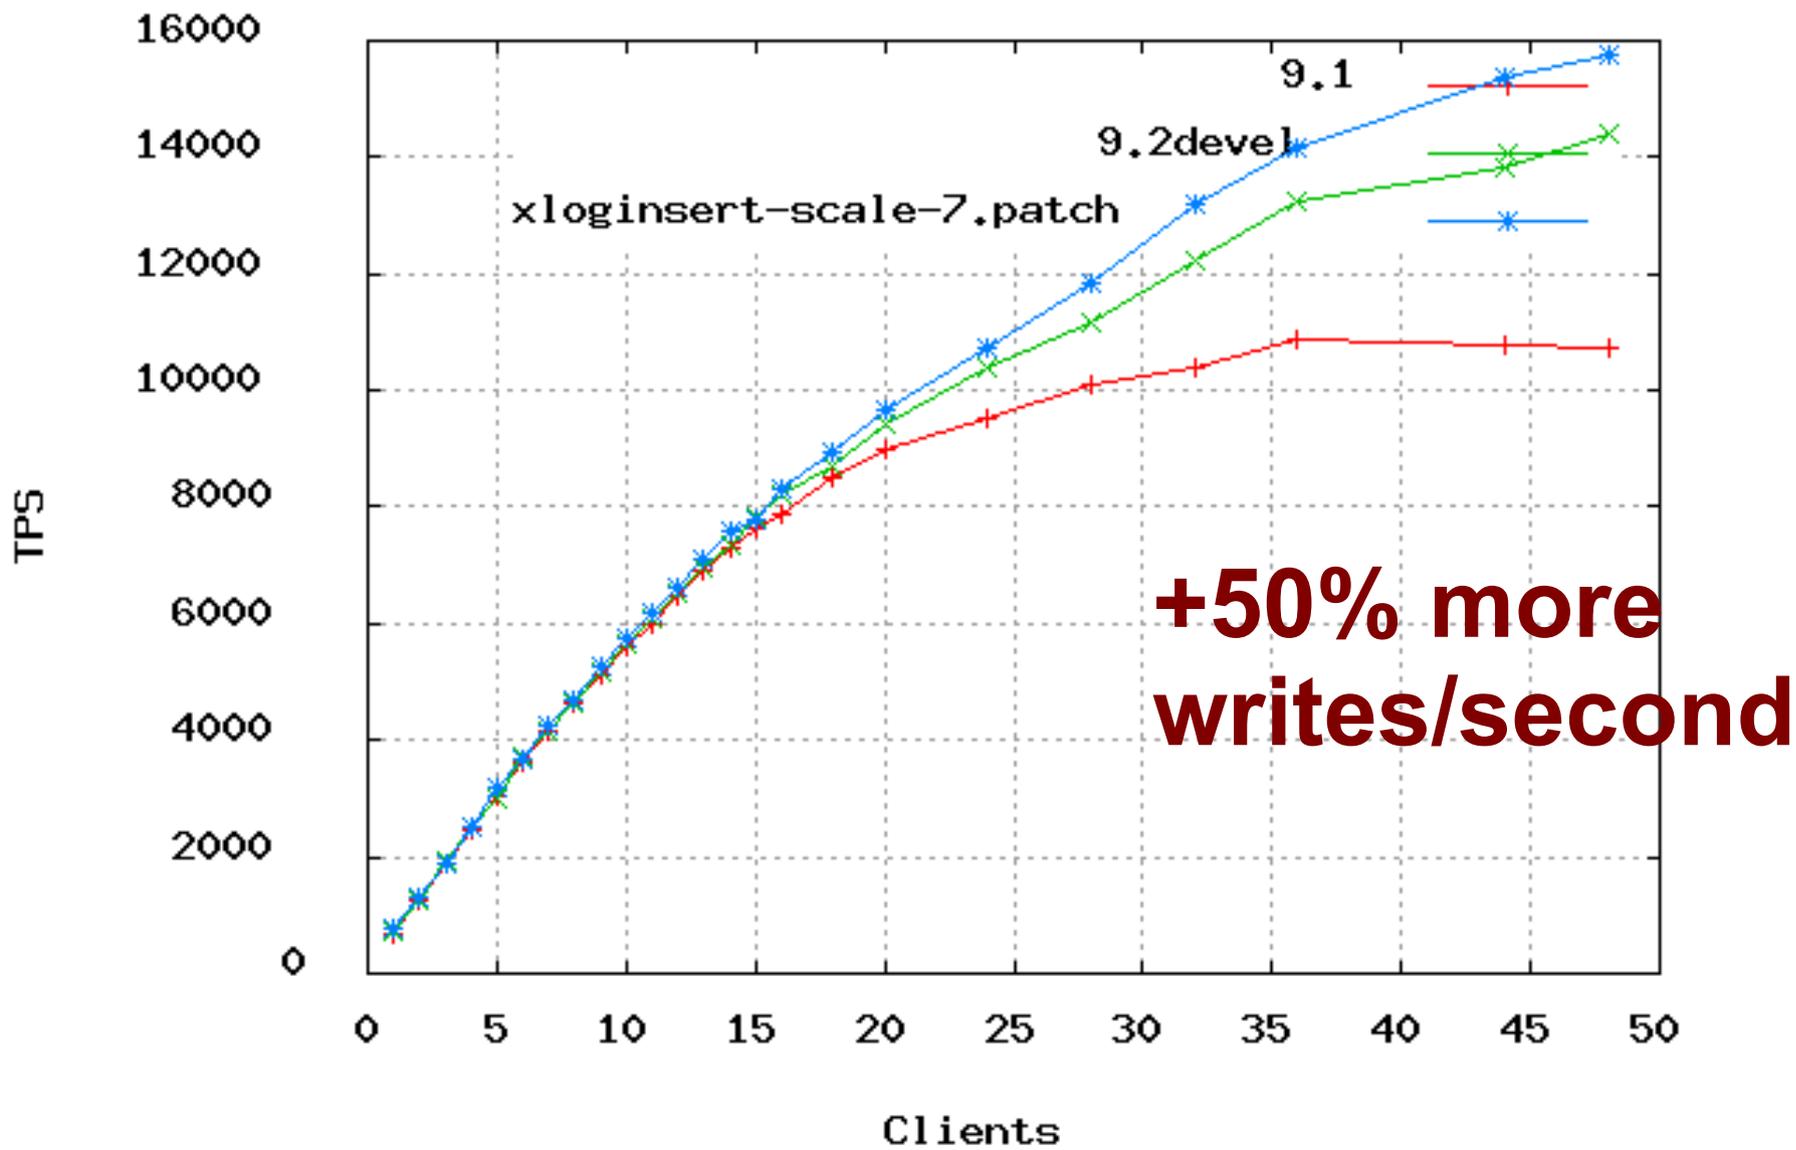
350,000 q.p.s.

Write Performance

- better group commit
- reduce WalInsertLock contention
- batch tuple copy
- reduced checkpoint activity

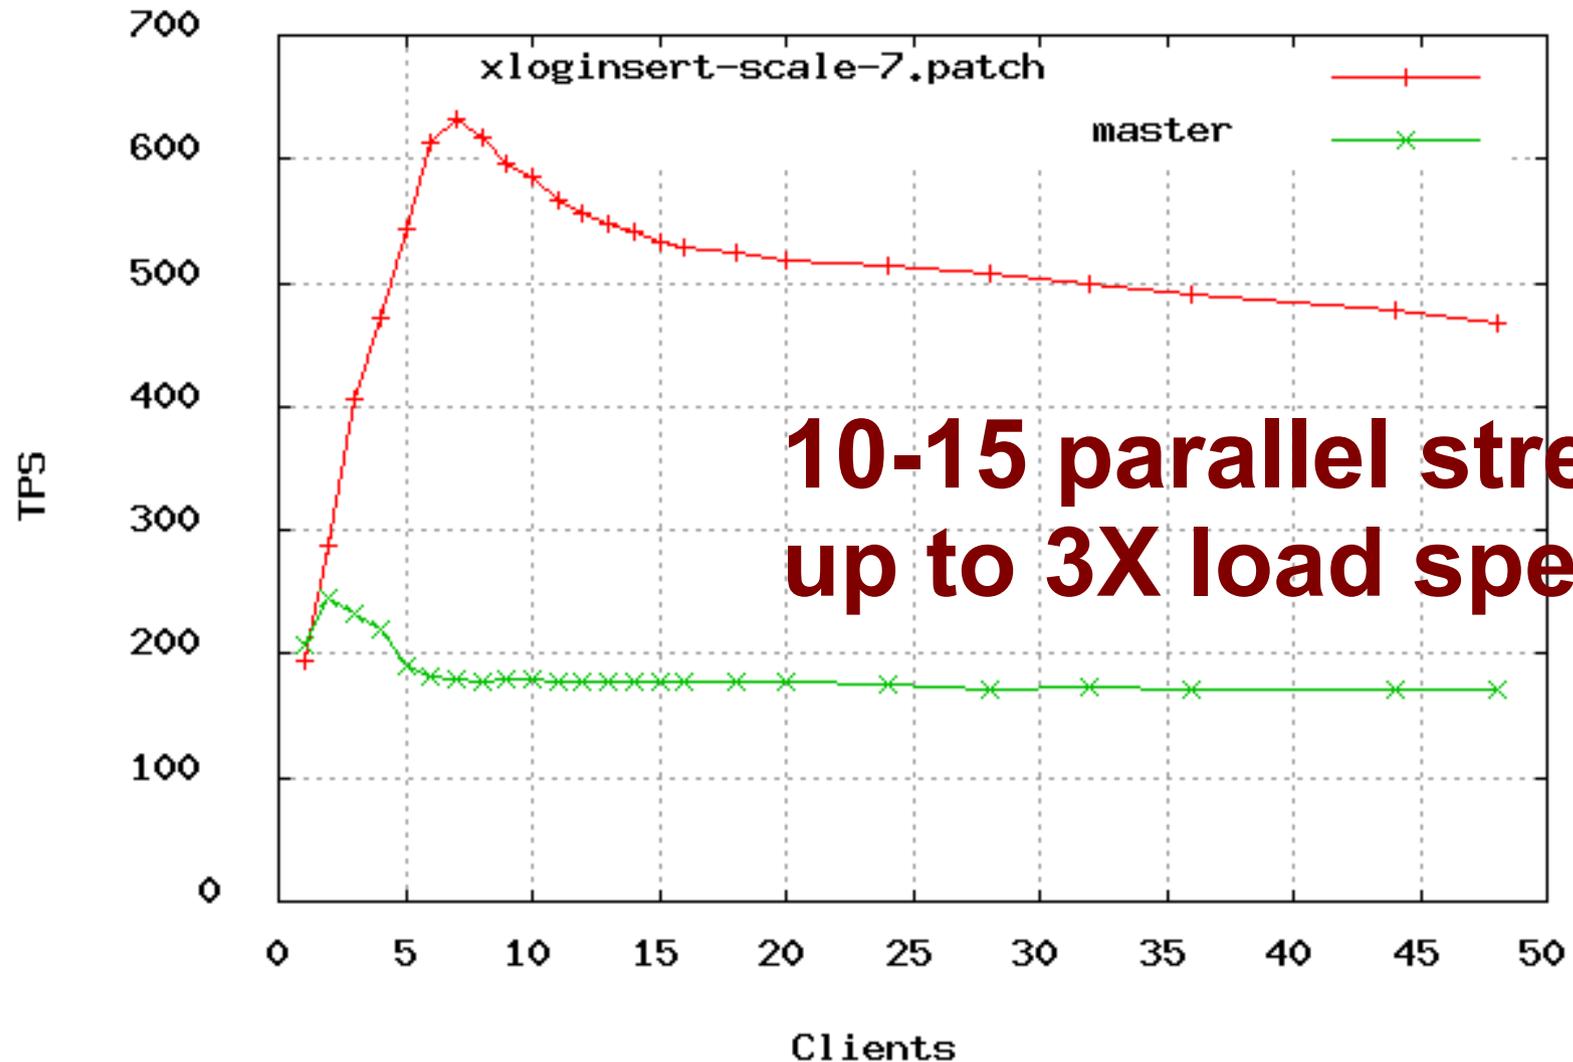
Write Scalability

pgbench transactions/sec (no branch update)



Parallel Bulk Load

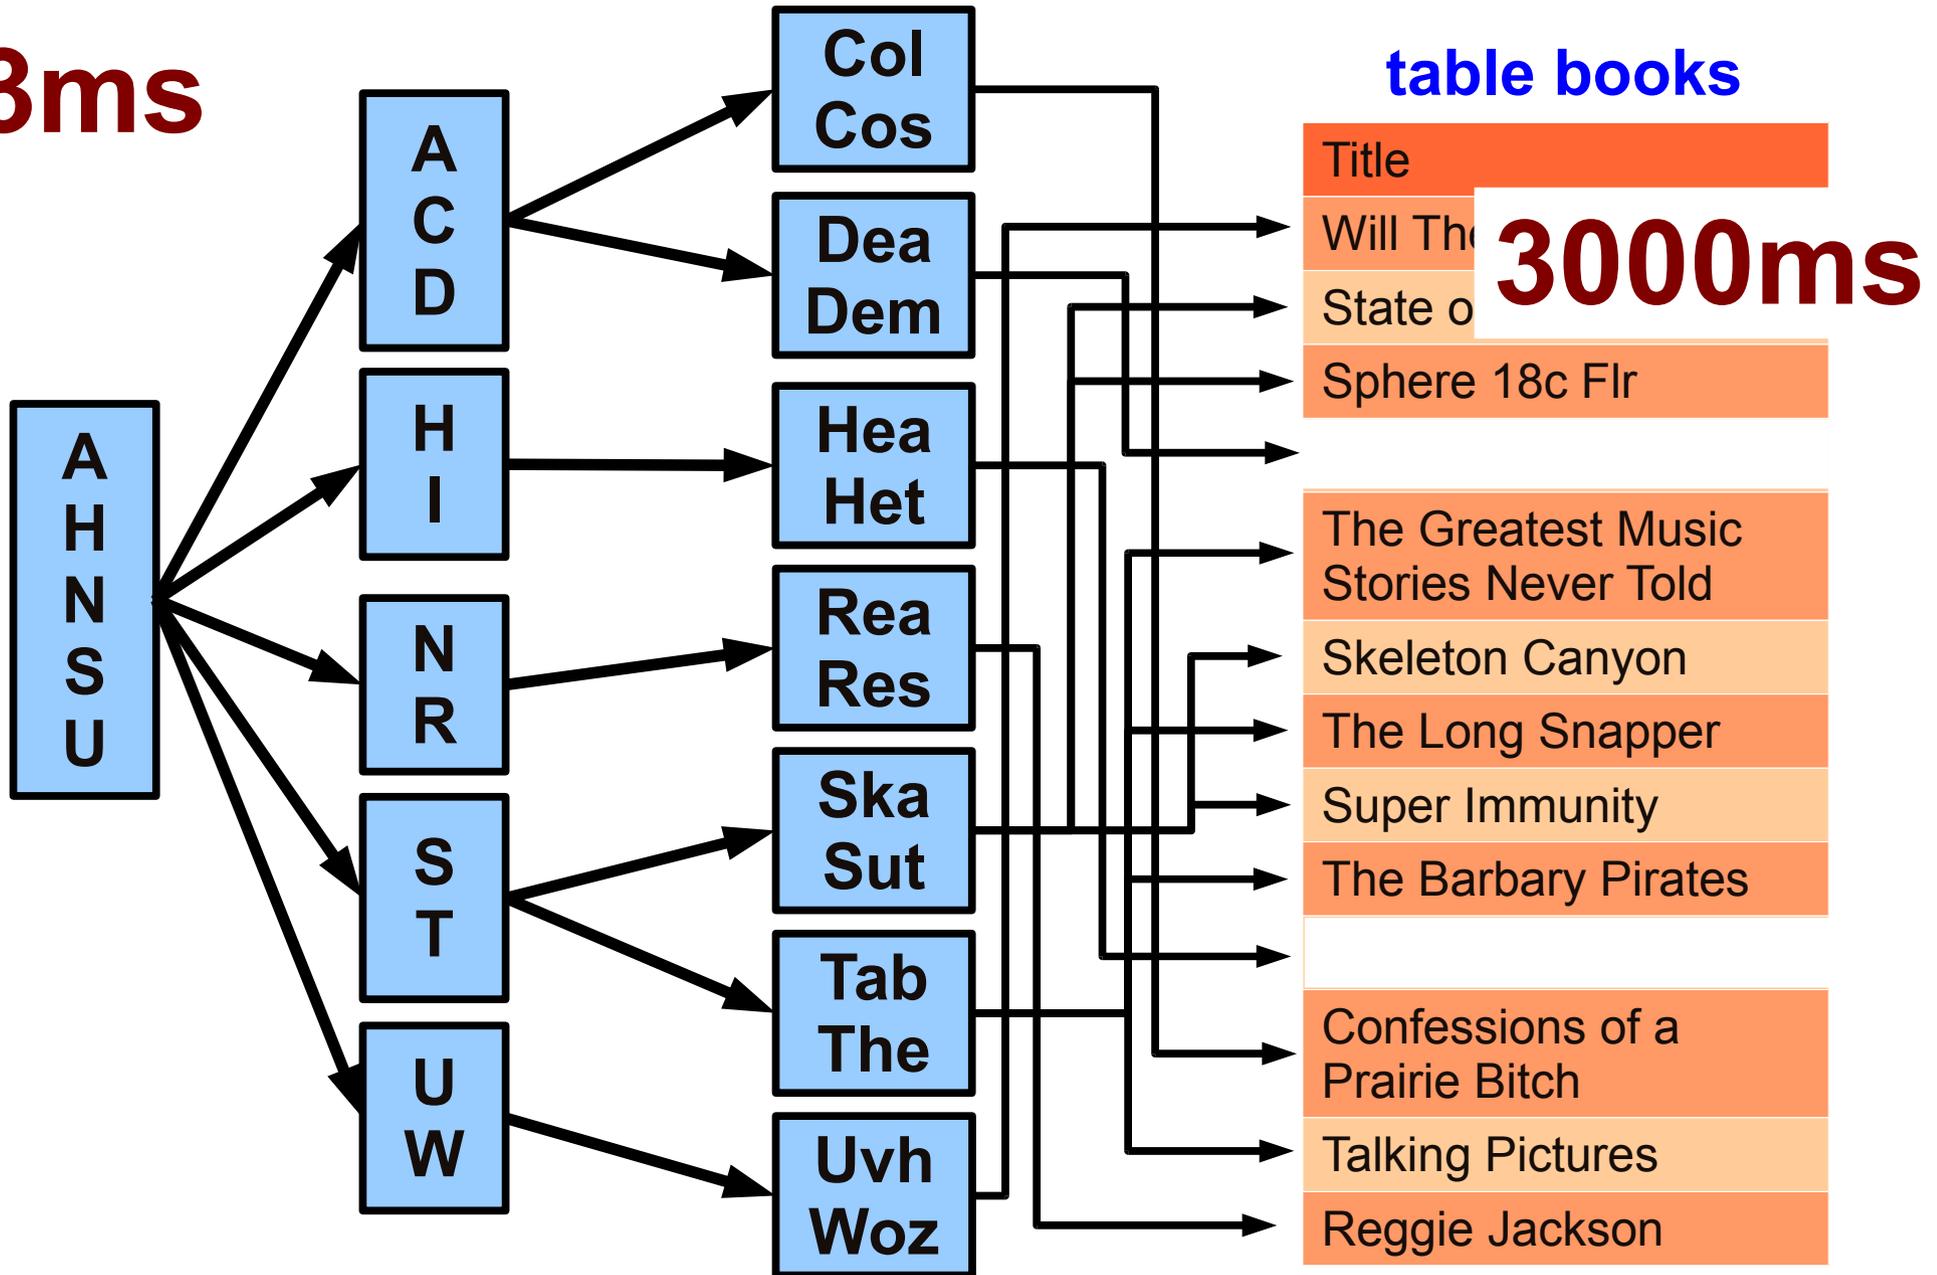
large 1000 row INSERT transactions/sec



**10-15 parallel streams
up to 3X load speed**

Index-Only Scans

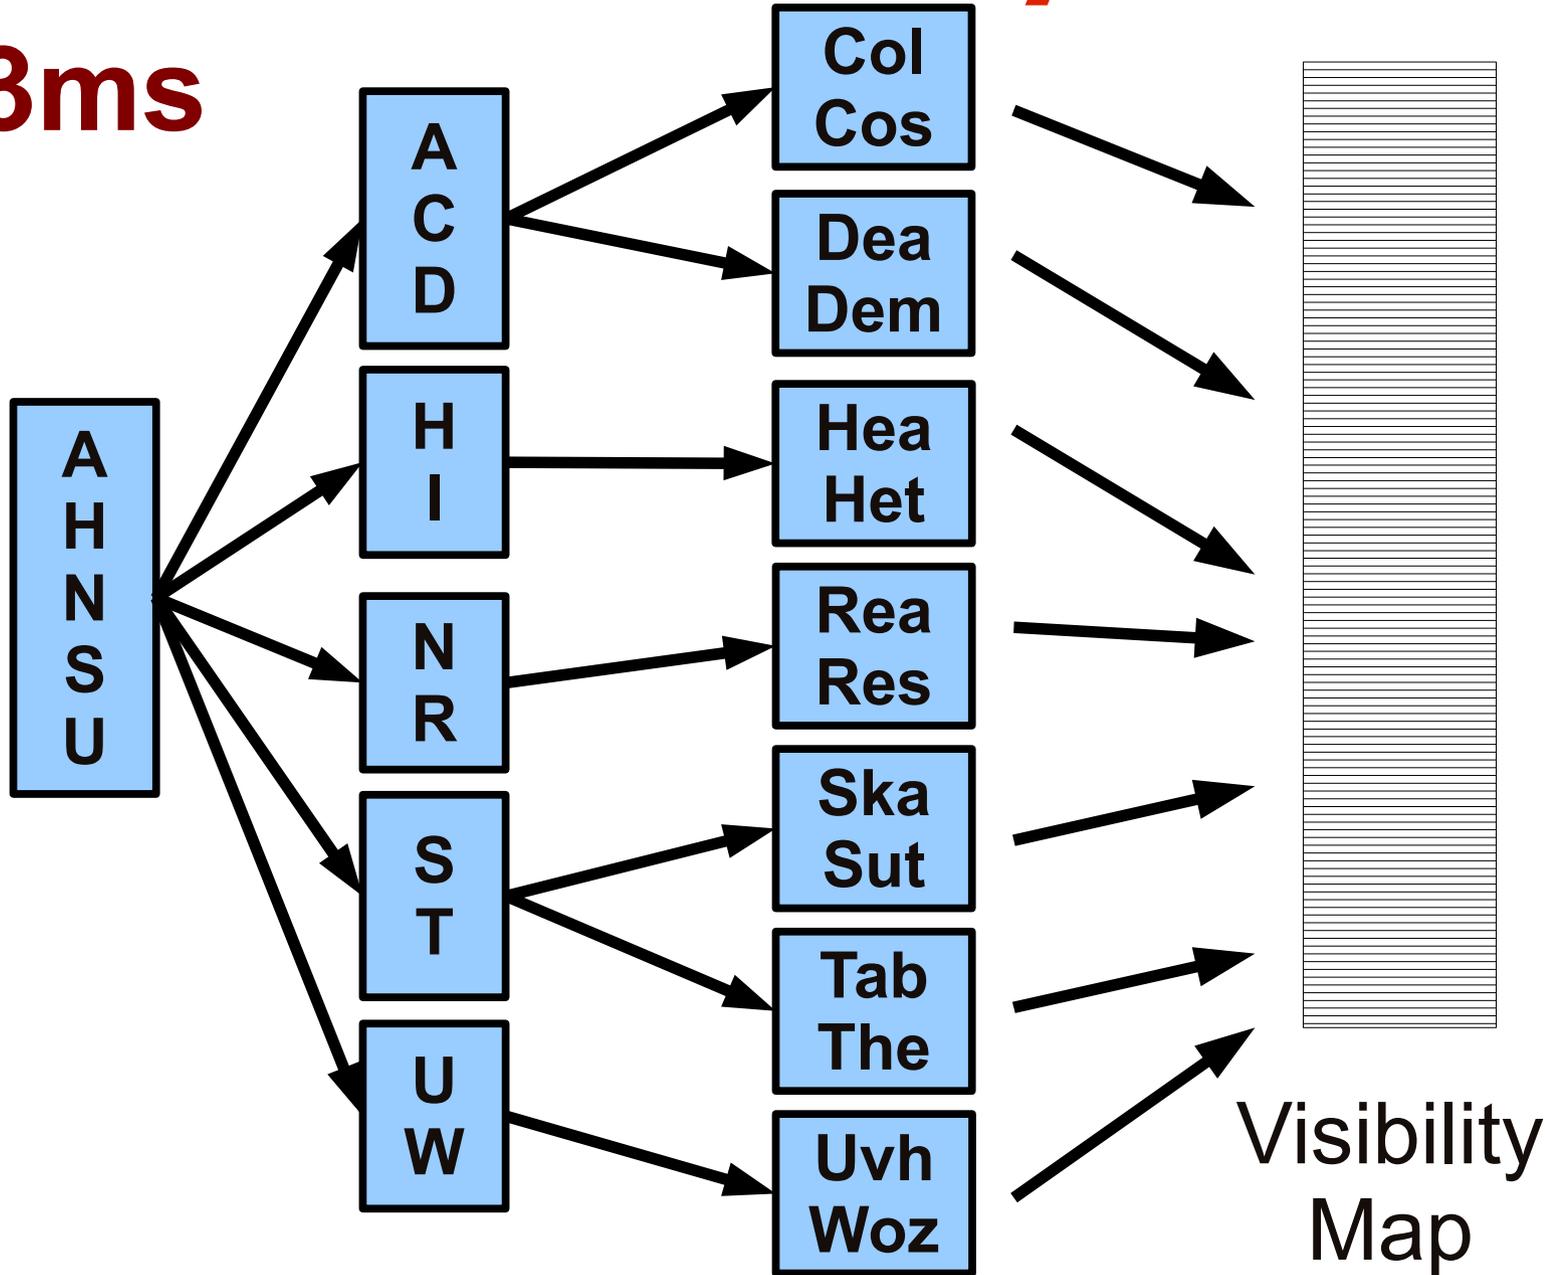
3ms



"count (*) is slow in
postgres"

Index-Only Scans

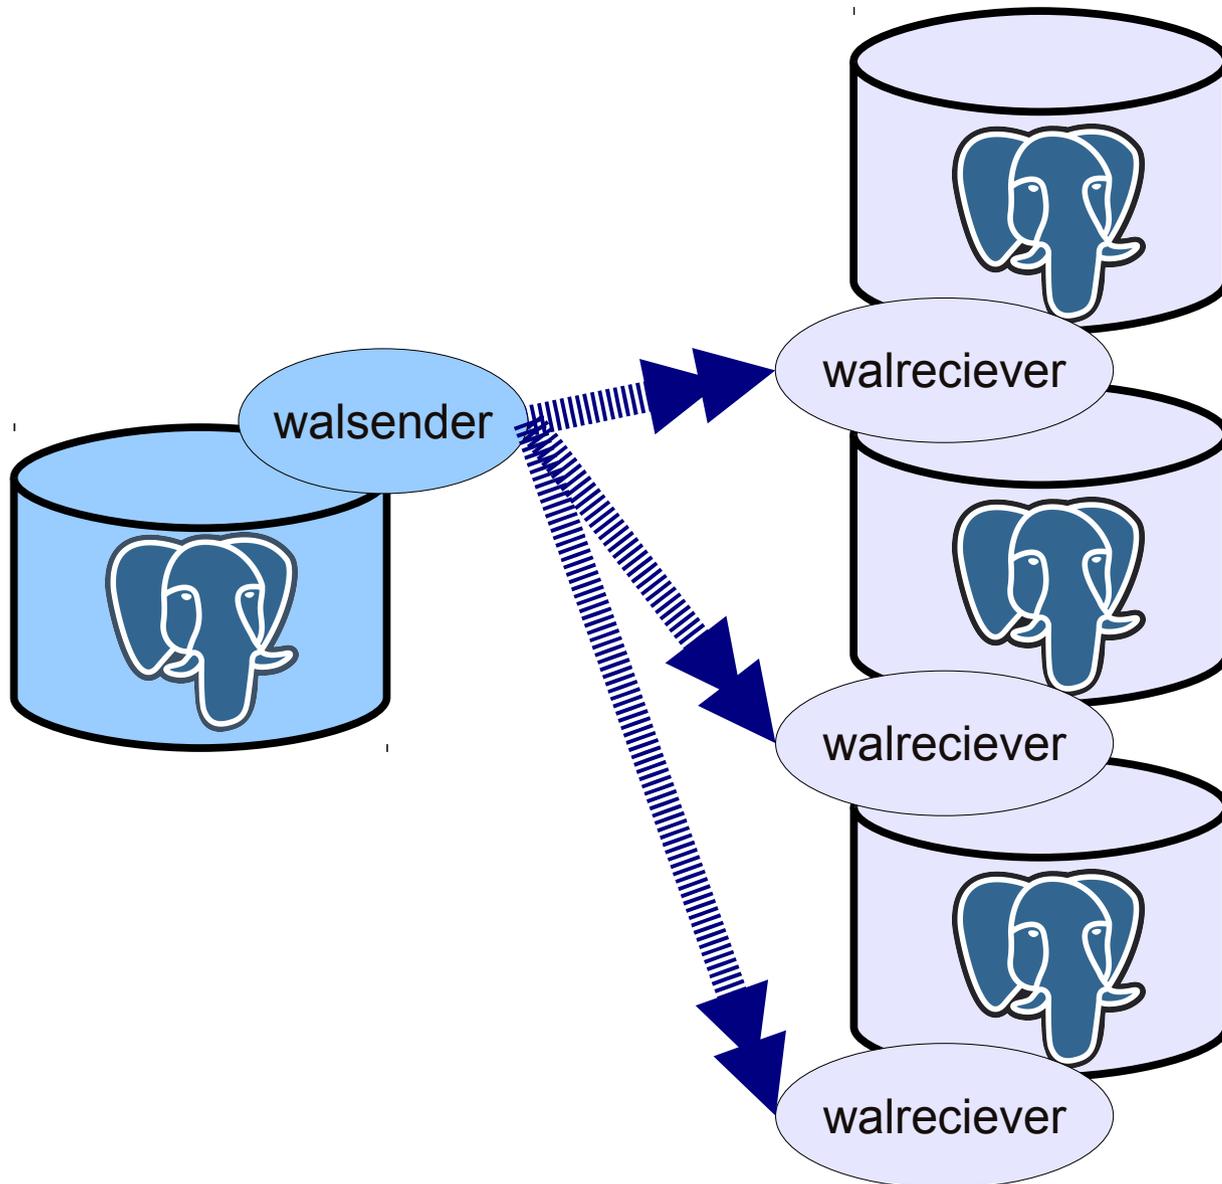
3ms



30ms

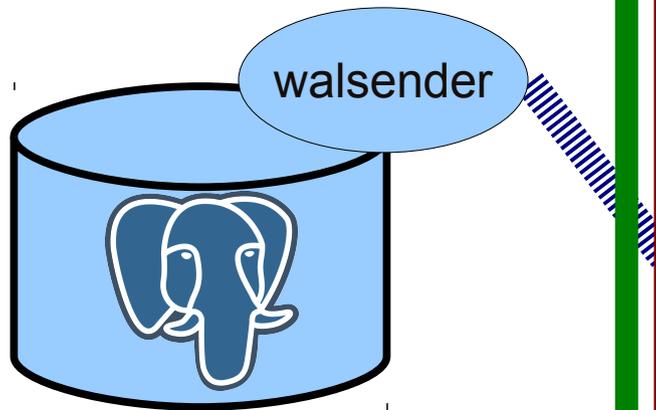
Visibility Map

Cascading Replication

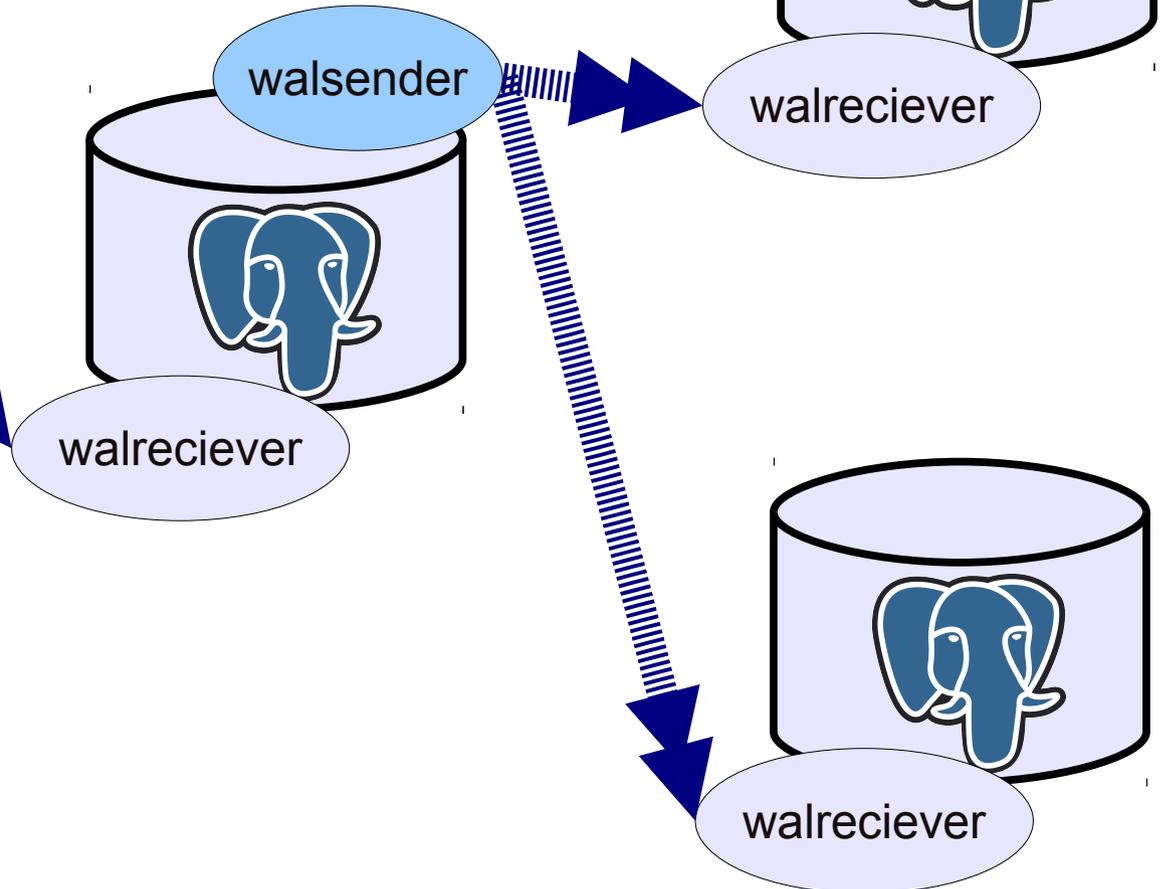


Cascading Replication

San Francisco



North Carolina

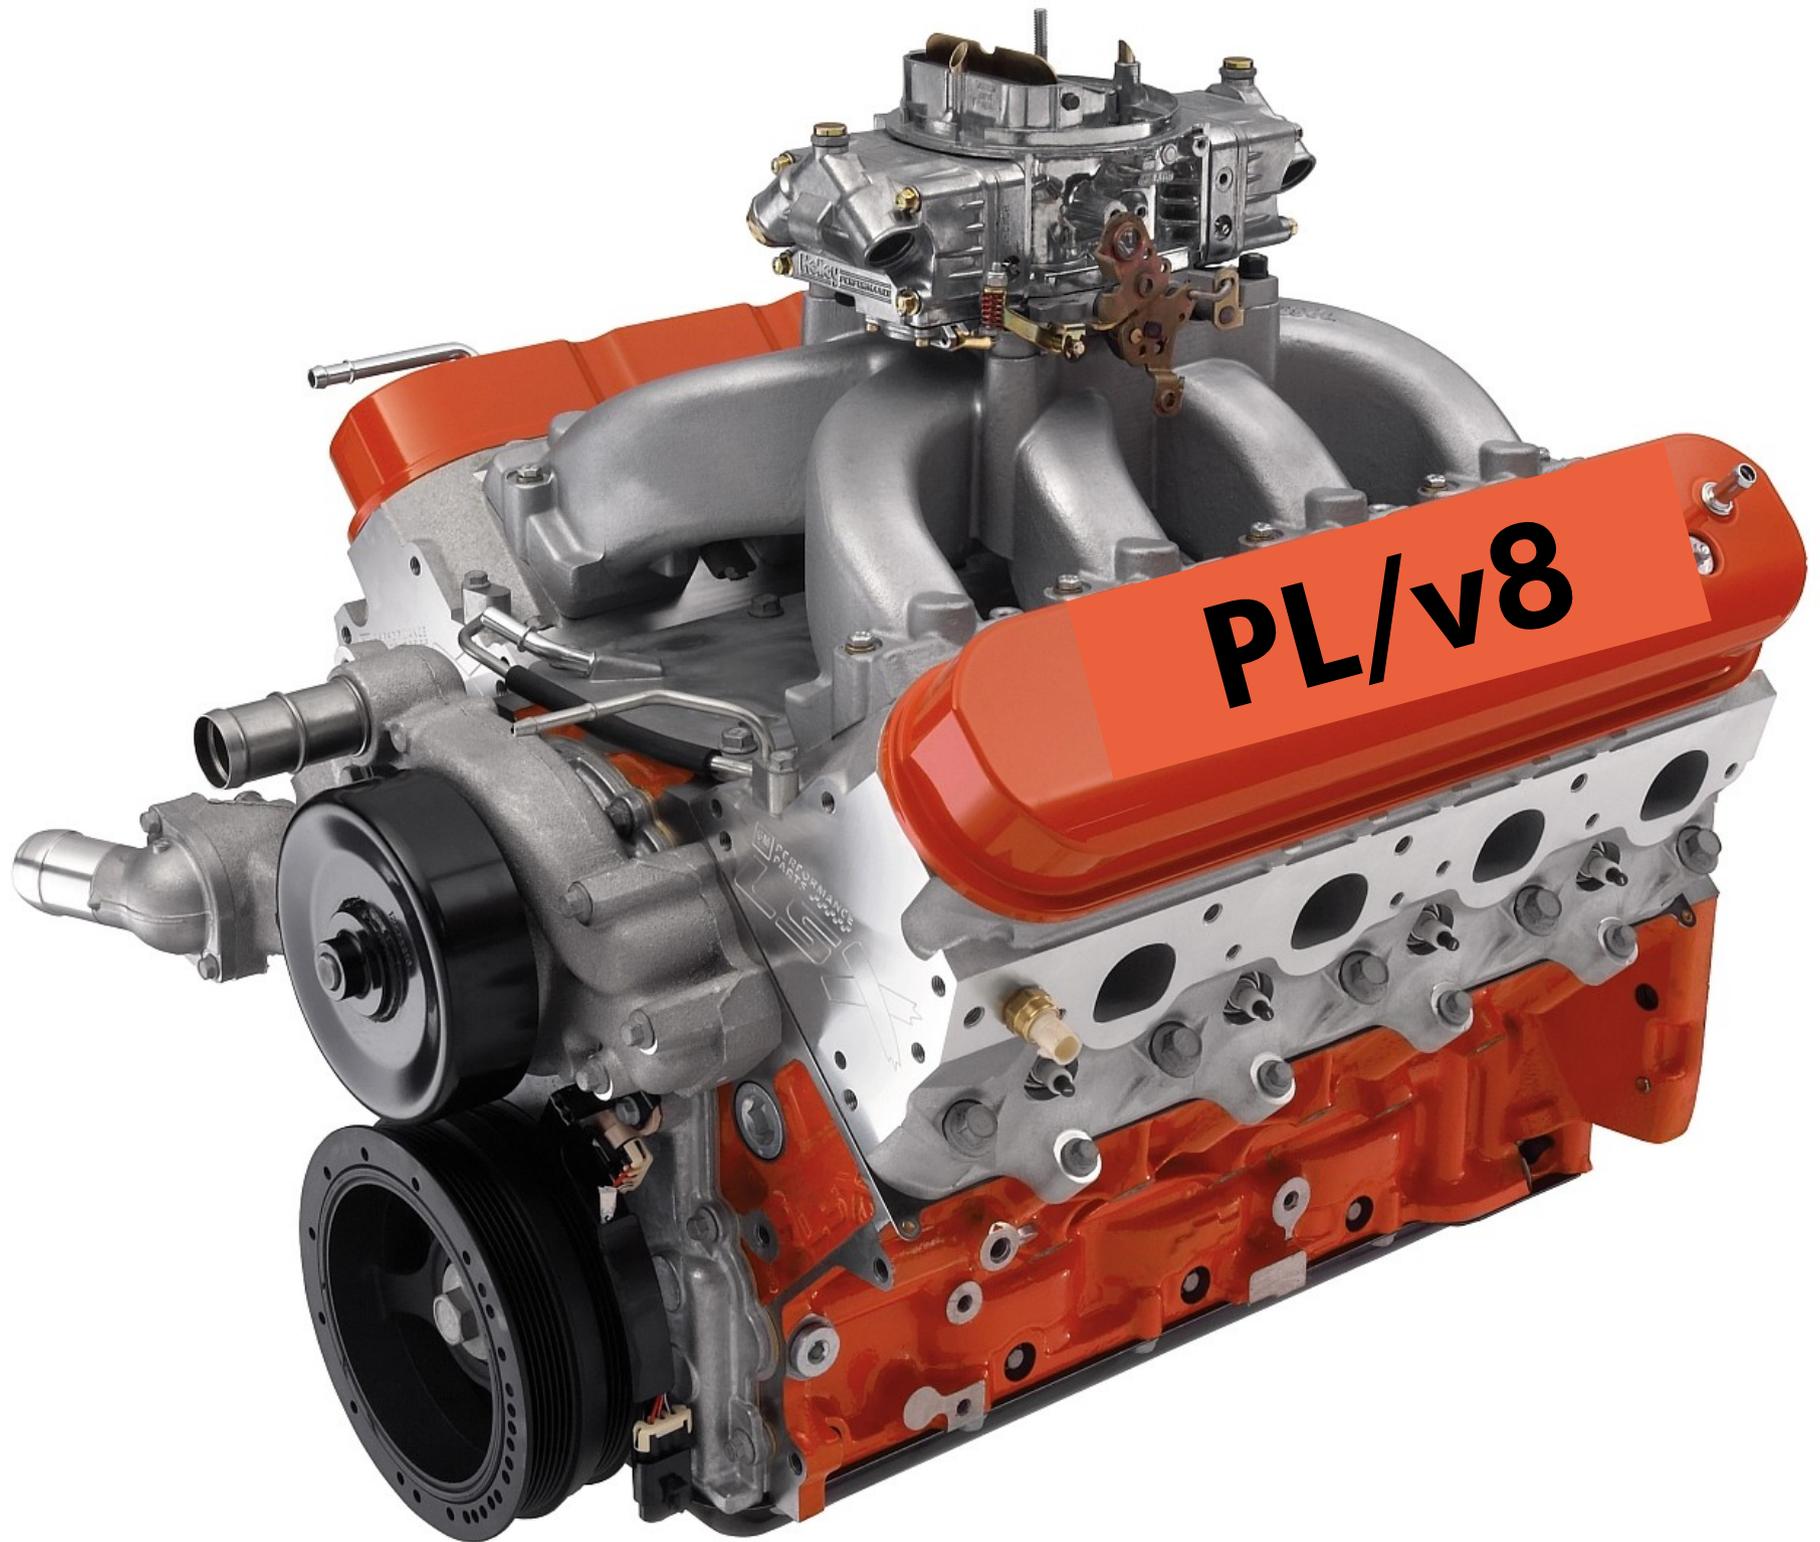


Replication Improvements

- pg_receivexlog
an xlog stream reader
- standby-only backup
no more snapshotting the master
- recv vs. write modes
for synchronous replication

JSON

- JSON data type
- array-to-json, row-to-json functions
- get query results as JSON!



PL/v8

PL/v8

- Based on Google's v8 JS engine
- write fast stored procedures in JavaScript
- create ad-hoc attribute indexes on JSON data

A data type and PL made for each other!

pl/coffee



PL/coffee

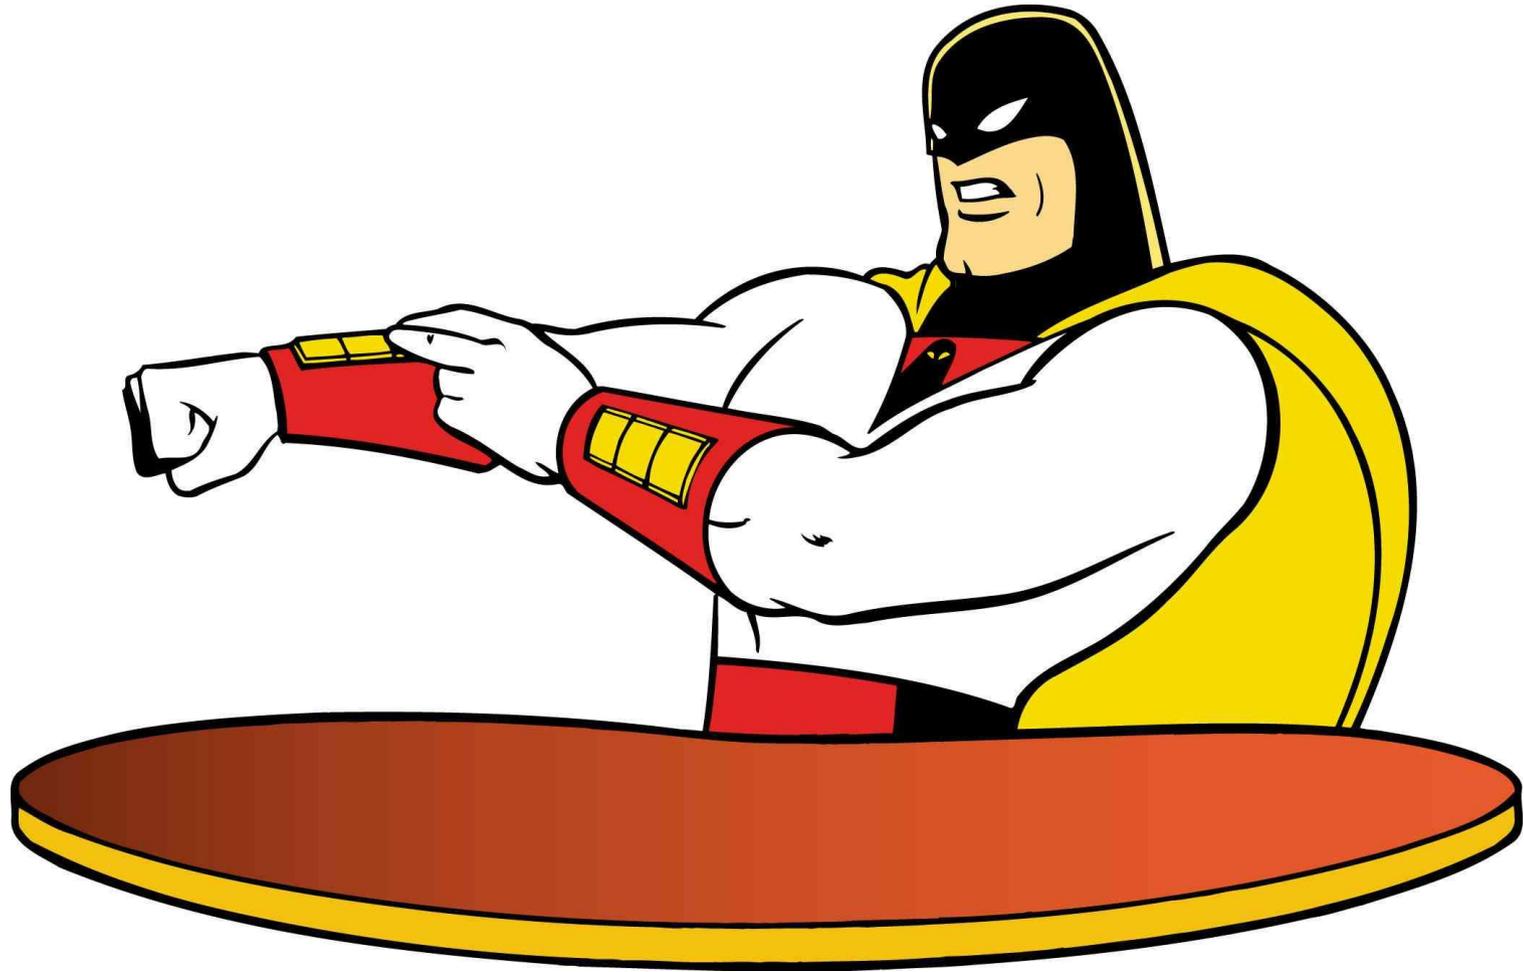
```
=# CREATE OR REPLACE FUNCTION
public.fibonacci(n integer)
  RETURNS integer LANGUAGE plcoffee
IMMUTABLE STRICT
AS $function$
fibonacci = (x) ->
  return 0 if x == 0
  return 1 if x == 1
  return fibonacci(x-1) + fibonacci(x-2)
return fibonacci n
$function$;
CREATE FUNCTION
```

Indexing

7.1: Generalized Search Tree
(GiST)

8.1: Generalized Inverted Index
(GIN)

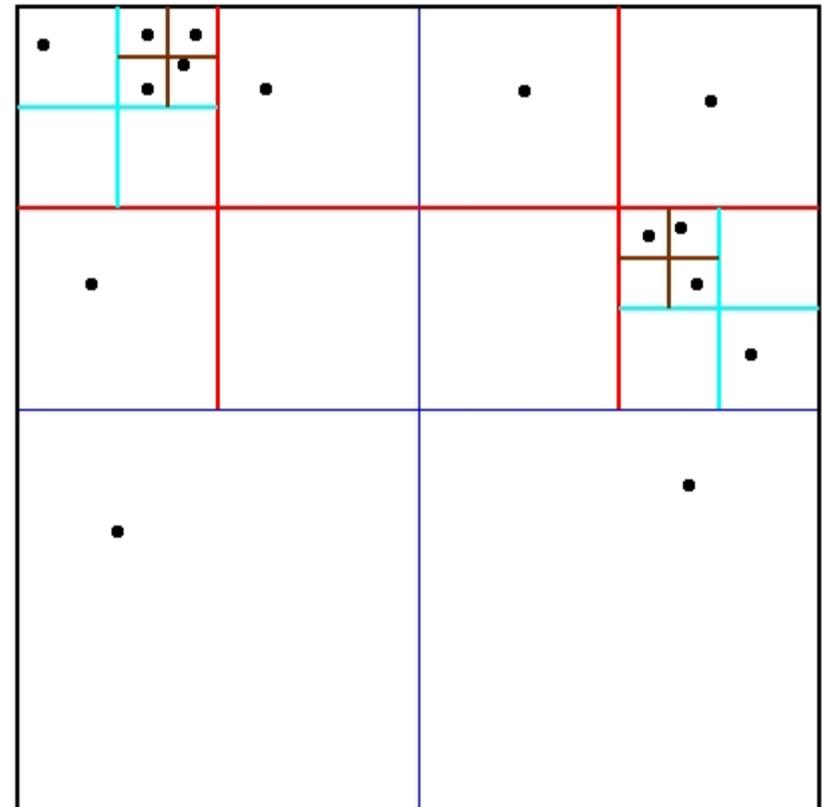
9.1: K-Nearest Neighbor
(KNN)



9.2: Space-GiST

SP-GiST

- Based on “Space-Partitioning Trees”
- Faster to read, faster to update than GiST
 - Will cause PostGIS 2.1!



SP-GiST

```
knn=# create index pt_gist_idx on geo  
using gist(point);
```

```
CREATE INDEX
```

```
Time: 36672.283 ms
```

```
Size: 153,124,864
```

```
knn=# create index pt_spgist_idx on geo  
using spgist(point);
```

```
CREATE INDEX
```

```
Time: 12805.530 ms ~ 3 times faster !
```

```
Size: 153,788,416 ~ the same size
```

Range Types

- Temporal range:
[2012-04-10, 2012-04-12)
- Alpha Index:
(Abbe, Babel]
- Linear distance:
(375.453, 374.441)



DDL pit stop

runtime DDL made easy

- DROP INDEX CONCURRENTLY
- NOT VALID CHECK constraints
- less rewriting for ALTER TABLE
- ALTER IF EXISTS
- RENAME domains, FDWs
- vacuum skips locked pages



Instrumentation

better monitoring

- new pg_stat_statements
- more autovacuum logging
- track_io_timing
- deadlock, temp file counters
- checkpoint timing in pg_stat_bgwriter

better EXPLAIN

- track buffers dirtied/written
- report filtered-out rows
- no-TIMING option

W-w-w-iiiiiiipeout!



didn't make it

- FK locks optimization
- data page checksums
- regexp indexing
- command triggers
- pgsql_fdw data federation
- parallel pg_dump

Other Features

- lower CPU wakeups (power saving)
- more sepgsql security
- optimizer improvements
- pg_hba.conf fixes
- XML improvements
- exportable snapshots

Other Features

- more pg_dump/pg_restore options
- more psql features
- more pg_upgrade fixes, improvements
- dropped support for antiques
 - dgux, nextstep, sunos4, svr4, ultrix4, univel, bsd

SÃO PAULO

Finish Line!



More Information

- Josh Berkus: josh@pgexperts.com
 - www.databasesoup.com
- pgExperts: www.pgexperts.com
- Events
 - S.E.L.F., OpenDBCon: June 7-11, Charlotte
 - pgOpen: Chicago, Sept. 14-17

Text and graphics of this talk are licensed Creative Commons Share-Alike, except where otherwise noted. Many photos in the presentation are used with or without permission and may not be reproduced or redistributed without permission of their owners. Graphs courtesy Robert Haas and Heikki Linnakangas. SP-Gist examples from Oleg Bartunov.

