

Postgres Interface Performance

Joe Conway

joe.conway@credativ.com

mail@joeconway.com

credativ International

October 22, 2014



Motivation

- Answer this question:
 - What is relative performance of ODBC, JDBC, and libpq?
- Specifically:
 - Connection speed
 - Query result retrieval
 - Alternate retrieval methods
 - Materializing to file
- Secondly:
 - Connection per query versus persistent
 - Set versus singleton operations
 - Local versus remote


Approach

- Seed database
- Launch command line clients
- Vary conditions
- Collect durations
- Analyze results

Major Components

- Command line clients
- PG Extension
- Driver, load, and result objects
- Driver function
- Launcher bash script

Flow



Clients

- Individual command line clients
 - Native language for interface
 - libpq - C
 - ODBC - C
 - JDBC - Java
 - As similar as possible source code structure
 - Identical command line arguments
 - Handle variety of test conditions

Extension

- C extension for command execution
 - Simple way to execute commands
 - Complete control over string escaping, etc.
 - Allow entire test to be table driven
 - Easy test result (duration) data collection

Driver Tables

- interfaces \Rightarrow libpq, ODBC, JDBC
- alternateclientcopy
 - \Rightarrow acc-n: normal interface
 - \Rightarrow acc-a: alternate interface
 - \Rightarrow acc-y: copy interface
- conninfos \Rightarrow lcl:ipc: local connection, BSD socket
 - \Rightarrow lcl:tcp: local connection, TCP/IP
 - \Rightarrow rmt:tcp: remote connection, TCP/IP
- printfracdest
 - \Rightarrow p-none:null: No print to file
 - \Rightarrow p-all:null: Print to /dev/null
 - \Rightarrow p-all:file: Print to file

Driver Tables

- sqlstratiter
 - ⇒ sql001:0:50k: 1 row, connect each time, run 50k times
 - ⇒ sql50k:0:001: 50k rows, connect each time, run 1 time
 - ⇒ sql10m:0:001: 10m rows, connect each time, run 1 time
 - ⇒ sql001:1:50k: 1 row, connect once, run 50k times
 - ⇒ sql50k:1:001: 50k rows, connect once, run 1 time
 - ⇒ sql10m:1:001: 10m rows, connect once, run 1 time

Driver View

- testcases
 - Generates all testcases from driver tables
 - Uses both JOIN (inner) and CROSS JOIN (Cartesian product)

```
CREATE OR REPLACE VIEW testcases AS
```

```
...
```

```
FROM interfaces i
```

```
JOIN conninfos c USING (ifname)
```

```
JOIN alternatclientcopy a USING (ifname)
```

```
CROSS JOIN printfracdest
```

```
CROSS JOIN sqlstratiter
```

```
...;
```

```
SELECT * FROM testcases LIMIT 1;
```

```
-[ RECORD 1 ]-----
```

```
testcase | ODBC:lcl:ipc:sql001:0:50k:acc-n:p-all:null
```

```
testcmd  | odbctest -c DSN=pgsql_ipc \
```

```
          -s "select * from t50k where x1 = 4242" \
```

```
          -p 1 -k 0 -n 50000 > /dev/null 2>&1
```



Load Tables

- t50k - 50 thousand rows

```

INSERT INTO t50k
SELECT
  id
  , '1aaaaaaaaaaaaaaaa[...len = 99...]aaaaaaaaaaaaaaaaaaaaaa'
  , '2bbbbbbbbbbbbbbbb[...len = 99...]bbbbbbbbbbbbbbbbbbbb'
  , '3cccccccccccccccc[...len = 99...]cccccccccccccccccccc'
  , '4aaaaaaaaaaaaaaaa[...len = 99...]aaaaaaaaaaaaaaaaaaaaaa'
  , '5bbbbbbbbbbbbbbbb[...len = 99...]bbbbbbbbbbbbbbbbbbbb'
  , '6cccccccccccccccc[...len = 99...]cccccccccccccccccccc'
  , '7aaaaaaaaaaaaaaaa[...len = 99...]aaaaaaaaaaaaaaaaaaaaaa'
  , '8bbbbbbbbbbbbbbbb[...len = 99...]bbbbbbbbbbbbbbbbbbbb'
  , '9cccccccccccccccc[...len = 99...]cccccccccccccccccccc'
FROM generate_series(1,50000) as t(id);
  
```

Load Tables

- t10m - 10 million rows

```

INSERT INTO t10m
SELECT
  id
  , '1aaaaaaaaaaaaa'
  , '2bbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbb'
  , '3cccccccccccccc'
  , '4aaaaaaaaaaaaaa'
  , '5bbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbb'
  , '6cccccccccccccc'
  , '7aaaaaaaaaaaaaa'
  , '8bbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbbb'
  , '9cccccccccccccc'
FROM generate_series(1,10000000) as t(id);
  
```

Results Table

- testresults - aggregate test case representation

```
SELECT * FROM testresults LIMIT 1;
-[ RECORD 1 ]-----
testcase | ODBC:lcl:ipc:sql001:0:50k:acc-n:p-all:null
iter     | 1
ts       | 2014-09-06 17:27:11.585011-07
duration | 238.23702
```

Results View

- vtestresults - expanded test case representation

```
SELECT * FROM vtestresults LIMIT 1;
-[ RECORD 1 ]-----
iface      | JDBC
loc        | lcl
comm       | tcp
sql        | sql001
connstrat  | 0
numiter    | 50k
extra      | acc-n
prnt       | p-all
dest       | file
iter       | 3
ts         | 2014-09-06 17:27:11.585011-07
duration   | 206.280236
```

Driver Function

- `executestcases()`
- Arguments
 - `path` - client binary location (command line clients)
 - `ldlibpath` - linked library location (`libpq.so`)
 - `classpath` - Java class location (JDBC and getopt jar files)
 - `outeriter` - how many times to execute all test cases
- Returns SETOF RECORD
 - `iter` - current outeriter
 - `tcase` - aggregate form test case identifier
 - `dur` - duration of test case run

Launcher script

- Kicks off the entire automated test
 - Sets variables for `executestcases()` arguments
 - Calls `executestcases()` using `psql`
 - Currently executes 288 test cases per run
 - Runs for many hours and produces many GB output files

Interface Variations

```
SELECT * FROM interfaces;
  ifname | testprog
-----+-----
  ODBC   | odbctest
  libpq  | libpqtest
  JDBC   | java jdbctest
(3 rows)
```

Connection Type Variations

```
SELECT * FROM conninfos;
```

conndesc	ifname	conninfo
lcl:ipc	ODBC	-c DSN=pgsql_ipc
lcl:tcp	ODBC	-c DSN=pgsql_lcl
rmt:tcp	ODBC	-c DSN=pgsql_rmt
lcl:ipc	libpq	-c postgresql://%2Ftmp/iftest?user=postgres
lcl:tcp	libpq	-c postgresql://localhost/iftest?user=postgres
rmt:tcp	libpq	-c postgresql://192.168.1.3/iftest?user=postgres
lcl:tcp	JDBC	-c jdbc:postgresql://localhost/iftest?user=postgres
rmt:tcp	JDBC	-c jdbc:postgresql://192.168.1.3/iftest?user=postgres

(8 rows)

Result Size, Conn Count, Num Executions

```
SELECT sqlstratiterdesc as desc,
       sql,
       connstrat as strat,
       numiter
FROM sqlstratiter;
```

desc	sql	strat	numiter
sql001:0:50k	-s "select * from t50k where id = 4242"	-k 0	-n 50000
sql50k:0:001	-s "select * from t50k"	-k 0	-n 1
sql10m:0:001	-s "select * from t10m"	-k 0	-n 1
sql001:1:50k	-s "select * from t50k where id = 4242"	-k 1	-n 50000
sql50k:1:001	-s "select * from t50k"	-k 1	-n 1
sql10m:1:001	-s "select * from t10m"	-k 1	-n 1

(6 rows)

Materialize (Y/N) and Destination

```
SELECT * FROM printfracdest;
```

pfddesc	printfrac	dest
p-all:null	-p 1	/dev/null
p-all:file	-p 1	/tmp/ifttest.out
p-none:null	-p 0	/dev/null

(3 rows)

Use Alternate Fetch or COPY Interface

```
SELECT * FROM alternateclientcopy;
```

```

accdesc | ifname | sw
-----+-----+-----
acc-n   | ODBC   |
acc-n   | libpq  |
acc-a   | libpq  | -a
acc-y   | libpq  | -y
acc-n   | JDBC   |
acc-y   | JDBC   | -y
(6 rows)
```

Results - Interface/Method

```

CREATE OR REPLACE VIEW ifmeth_cat_results AS
SELECT sql || '$$: $$ || numiter || '$$: $$ || loc || '$$: $$ || comm ||
       '$$: $$ || connstrat || '$$: $$ || prnt || '$$: $$ || dest AS testcase
,iface || '$$: $$ || extra AS candidate
,round(duration::numeric, 5) AS duration
,round(drang::numeric, 5) AS drange, grpsz
,round(((duration - min(duration) OVER w) /
        (max(duration) OVER w - min(duration) OVER w)
        )::numeric, 5) AS nml_dur
,round(((drange - min(drang) OVER w) /
        (max(drang) OVER w - min(drang) OVER w)
        )::numeric, 5) AS nml_drng
FROM (SELECT iface,loc,comm,connstrat,sql,numiter,extra,prnt,dest
       ,count(1) as grpsz, avg(duration) AS duration
       ,max(duration) - min(duration) AS drange
       FROM vtestresults GROUP BY 1,2,3,4,5,6,7,8,9) as t
WINDOW w AS (PARTITION BY sql,numiter,loc,comm,connstrat,prnt,dest
              ORDER BY duration
              ROWS BETWEEN UNBOUNDED PRECEDING AND UNBOUNDED FOLLOWING);

```

Results - Interface/Method Summarized

```
SELECT candidate
      ,round(AVG(grpsz),1) AS avg_grpsz
      ,round(AVG(nml_dur),3) AS avg_nml_dur
      ,round(AVG(nml_drng),3) AS avg_nml_drng
FROM ifmeth_cat_results
GROUP BY candidate
ORDER BY 3,4,1;
```

candidate	avg_grpsz	avg_nml_dur	avg_nml_drng
libpq:acc-a	9.0	0.008	0.285
libpq:acc-n	9.0	0.055	0.224
libpq:acc-y	9.0	0.097	0.416
JDBC:acc-y	9.0	0.231	0.279
ODBC:acc-n	9.0	0.748	0.523
JDBC:acc-n	9.0	0.851	0.600

(6 rows)

Results - Location/Transport

```

CREATE OR REPLACE VIEW loccomm_cat_results AS
SELECT iface || '$': '$' || extra || '$': '$' || sql || '$': '$' || numiter ||
    '$': '$' || connstrat || '$': '$' || prnt || '$': '$' || dest AS testcase
,loc || '$': '$' || comm AS candidate
,round(duration::numeric, 5) AS duration
,round(drange::numeric, 5) AS drange, grpsz
,round(((duration - min(duration) OVER w) /
    (max(duration) OVER w - min(duration) OVER w)
)::numeric, 5) AS nml_dur
,round(((drange - min(drange) OVER w) /
    (max(drange) OVER w - min(drange) OVER w)
)::numeric, 5) AS nml_drng
FROM (SELECT iface,loc,comm,connstrat,sql,numiter,extra,prnt,dest
    ,count(1) as grpsz, avg(duration) AS duration
    ,max(duration) - min(duration) AS drange
    FROM vtestresults GROUP BY 1,2,3,4,5,6,7,8,9) as t
WINDOW w AS (PARTITION BY sql,numiter,iface,extra,connstrat,prnt,dest
    ORDER BY duration
    ROWS BETWEEN UNBOUNDED PRECEDING AND UNBOUNDED FOLLOWING);

```


Results - Location/Transport Summarized

```
SELECT candidate
      ,round(AVG(grpsz),1) AS avg_grpsz
      ,round(AVG(nml_dur),3) AS avg_nml_dur
      ,round(AVG(nml_drng),3) AS avg_nml_drng
FROM loccomm_cat_results
GROUP BY candidate
ORDER BY 3,4,1;
```

candidate	avg_grpsz	avg_nml_dur	avg_nml_drng
lcl:ipc	9.0	0.001	0.398
lcl:tcp	9.0	0.108	0.484
rmt:tcp	9.0	1.000	0.538

(3 rows)

Results - Query/Connect Strategy Summarized

```
SELECT testcase, round(avg(duration),3) AS avg_dur
FROM ifmeth_cat_results
WHERE testcase IN (
    'sql50k:001:lcl:ipc:0:p-all:file'
  , 'sql50k:001:lcl:ipc:1:p-all:file'
  , 'sql50k:001:lcl:tcp:0:p-all:file'
  , 'sql50k:001:lcl:tcp:1:p-all:file'
  , 'sql50k:001:rmt:tcp:0:p-all:file'
  , 'sql50k:001:rmt:tcp:1:p-all:file'
  , 'sql001:50k:lcl:ipc:0:p-all:file'
  , 'sql001:50k:lcl:ipc:1:p-all:file'
  , 'sql001:50k:lcl:tcp:0:p-all:file'
  , 'sql001:50k:lcl:tcp:1:p-all:file'
  , 'sql001:50k:rmt:tcp:0:p-all:file'
  , 'sql001:50k:rmt:tcp:1:p-all:file'
)
GROUP BY testcase
ORDER BY 2;
```

Results - Query/Connect Strategy Summarized

testcase	avg_dur
sql50k:001:lcl:ipc:0:p-all:file	0.293
sql50k:001:lcl:ipc:1:p-all:file	0.299
sql50k:001:lcl:tcp:0:p-all:file	0.758
sql50k:001:lcl:tcp:1:p-all:file	0.784
sql50k:001:rmt:tcp:0:p-all:file	1.036
sql50k:001:rmt:tcp:1:p-all:file	1.038
sql001:50k:lcl:ipc:1:p-all:file	10.245
sql001:50k:lcl:tcp:1:p-all:file	15.586
sql001:50k:rmt:tcp:1:p-all:file	40.944
sql001:50k:lcl:ipc:0:p-all:file	186.176
sql001:50k:lcl:tcp:0:p-all:file	203.528
sql001:50k:rmt:tcp:0:p-all:file	428.544

(12 rows)

sql001:50k:lcl:ipc:0:p-all:file

- sql001:50k ⇒ Query 1 row, 50k times
- lcl:ipc ⇒ Local host via IPC connection
- 0 ⇒ Connect every query
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql001:50k:lcl:ipc:0:p-all:file';
```

candidate	duration
libpq:acc-n	167.08301
libpq:acc-a	168.08354
libpq:acc-y	169.53112
ODBC:acc-n	240.00766

(4 rows)

sql001:50k:lcl:ipc:1:p-all:file

- sql001:50k ⇒ Query 1 row, 50k times
- lcl:ipc ⇒ Local host via IPC connection
- 1 ⇒ Connect once
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql001:50k:lcl:ipc:1:p-all:file';
```

candidate	duration
libpq:acc-n	7.72111
libpq:acc-a	8.30801
libpq:acc-y	8.44191
ODBC:acc-n	16.50867

(4 rows)



sql001:50k:lcl:tcp:0:p-all:file

- sql001:50k ⇒ Query 1 row, 50k times
- lcl:tcp ⇒ Local host via TCP/IP connection
- 0 ⇒ Connect every query
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql001:50k:lcl:tcp:0:p-all:file';
```

```

candidate | duration
-----+-----
libpq:acc-n | 184.76682
libpq:acc-a | 185.10930
libpq:acc-y | 187.35046
JDBC:acc-y | 200.94066
JDBC:acc-n | 206.43863
ODBC:acc-n | 256.56118
(6 rows)
```

sql001:50k:lcl:tcp:1:p-all:file

- sql001:50k ⇒ Query 1 row, 50k times
- lcl:tcp ⇒ Local host via TCP/IP connection
- 1 ⇒ Connect once
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql001:50k:lcl:tcp:1:p-all:file';
```

candidate	duration
libpq:acc-a	11.40072
libpq:acc-n	11.62985
libpq:acc-y	12.67647
JDBC:acc-y	15.00848
JDBC:acc-n	20.46275
ODBC:acc-n	22.33883

(6 rows)

sql001:50k:rmt:tcp:0:p-all:file

- sql001:50k ⇒ Query 1 row, 50k times
- rmt:tcp ⇒ Remote host via TCP/IP connection
- 0 ⇒ Connect every query
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql001:50k:rmt:tcp:0:p-all:file';
```

candidate	duration
libpq:acc-n	380.87894
libpq:acc-a	381.20051
libpq:acc-y	398.77627
JDBC:acc-y	437.91512
JDBC:acc-n	443.36177
ODBC:acc-n	529.13435

(6 rows)

sql001:50k:rmt:tcp:1:p-all:file

- sql001:50k ⇒ Query 1 row, 50k times
- rmt:tcp ⇒ Remote host via TCP/IP connection
- 1 ⇒ Connect once
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql001:50k:rmt:tcp:1:p-all:file';
```

candidate	duration
libpq:acc-n	26.79044
libpq:acc-a	26.95195
libpq:acc-y	41.10257
JDBC:acc-y	44.27100
ODBC:acc-n	52.93944
JDBC:acc-n	53.60791

(6 rows)

sql50k:001:rmt:tcp:1:p-all:file

- sql50k:001 ⇒ Query 50k rows, 1 time
- rmt:tcp ⇒ Remote host via TCP/IP connection
- 1 ⇒ Connect once
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql50k:001:rmt:tcp:1:p-all:file';
```

candidate	duration
libpq:acc-a	0.44083
libpq:acc-y	0.55278
libpq:acc-n	0.60207
JDBC:acc-y	0.67728
ODBC:acc-n	0.94350
JDBC:acc-n	3.00985

(6 rows)

sql10m:001:rmt:tcp:1:p-all:file

- sql10m:001 ⇒ Query 10m rows, 1 time
- rmt:tcp ⇒ Remote host via TCP/IP connection
- 1 ⇒ Connect once
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql10m:001:rmt:tcp:1:p-all:file';
```

candidate	duration
libpq:acc-a	36.10610
libpq:acc-y	41.68094
JDBC:acc-y	41.85082
libpq:acc-n	52.39856
ODBC:acc-n	90.92070
JDBC:acc-n	416.97683

(6 rows)

sql10m:001:lcl:tcp:1:p-all:file

- sql10m:001 ⇒ Query 10m rows, 1 time
- lcl:tcp ⇒ Local host via TCP/IP connection
- 1 ⇒ Connect once
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql10m:001:lcl:tcp:1:p-all:file';
```

candidate	duration
libpq:acc-y	29.37412
libpq:acc-a	31.62445
JDBC:acc-y	31.87552
libpq:acc-n	40.85718
ODBC:acc-n	77.09062
JDBC:acc-n	396.85311

(6 rows)

sql10m:001:lcl:ipc:1:p-all:file

- sql10m:001 ⇒ Query 10m rows, 1 time
- lcl:ipc ⇒ Local host via IPC connection
- 1 ⇒ Connect once
- p-all:file ⇒ print every row to file

```
SELECT candidate, duration
FROM ifmeth_cat_results
WHERE testcase = 'sql10m:001:lcl:ipc:1:p-all:file';
```

```
  candidate | duration
-----+-----
 libpq:acc-y | 27.45592
 libpq:acc-a | 30.03265
 libpq:acc-n | 36.35336
 ODBC:acc-n  | 69.42425
(4 rows)
```



all:acc-n:sql50k:001:all:p-all:file

- all:acc-n ⇒ All interfaces, normal method
- sql50k:001 ⇒ 50k rows, one time
- p-all:file ⇒ print every row to file

```
SELECT testcase, candidate, duration
FROM loccomm_cat_results
WHERE testcase LIKE '%:acc-n:sql50k:001:1:p-all:file'
ORDER BY 1, 3, 2;
```

testcase	candidate	duration
libpq:acc-n:sql50k:001:1:p-all:file	lcl:ipc	0.21962
libpq:acc-n:sql50k:001:1:p-all:file	lcl:tcp	0.32501
libpq:acc-n:sql50k:001:1:p-all:file	rmt:tcp	0.60207
ODBC:acc-n:sql50k:001:1:p-all:file	lcl:ipc	0.56753
ODBC:acc-n:sql50k:001:1:p-all:file	lcl:tcp	0.70518
ODBC:acc-n:sql50k:001:1:p-all:file	rmt:tcp	0.94350
JDBC:acc-n:sql50k:001:1:p-all:file	lcl:tcp	2.72273
JDBC:acc-n:sql50k:001:1:p-all:file	rmt:tcp	3.00985

(8 rows)

Future Work

- Additional interfaces
 - Different versions of JDBC library
 - Alternate JDBC library
 - npgsql (.Net data provider)
 - py-postgresql (native python interface)
- Additional alternate fetch strategies
 - Prepared statements
 - Server-side COPY with client-side file fetch
- Cost of SSL connection

Questions?

Thank You!

joe.conway@credativ.com

mail@joeconway.com