

PgBench – Work in Progress

Fabien Coelho

MINES ParisTech, PSL Research University

PostgreSQL Session #9, Paris – November 17, 2017

1 PgBench

- History
- Capabilities
- Caveats

3 Future

- Needs
- CommitFest
- Conclusion

2 Performance

- Overheads
- Loading
- Connection Costs
- SSL Costs
- Index
- Fill Factor

PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion



PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future

Needs
CommitFest
Conclusion

Simple tool based on TPC-B

Tatsuo Ishii 2000

- external... then `contrib/`
- initialization and scale `-i -s 10`
- 2 benchmarks `-t 10000 -c 4`



Tatsuo Ishii by Oleg Bartunov

Visible *and invisible* developments

2001-

- Initializing
- Scripting
- Running
- Reporting
- Using...
- Debugging
- Refactoring
- Testing

PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion

<code>COPY</code>	initialization	Takahiro Itagaki	2007
<code>FILLFACTOR</code>	taux de remplissage	Pavan Deolasee	2007
<code>UNLOGGED</code>	tables	Robert Haas	2011
<code>TABLESPACE</code>	on tables or index	–	2011
<code>--foreign-key</code>	declarations	Jeff Janes	2012
<code>-I ...</code>	initialization steps	Masahiko Sawada	2017

PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion

<code>-N</code>	simple update	Tatsuo Ishii	2002
<code>-f</code>	script file	Tomoaki Sato	2005
<code>\set</code>	basic arithmetic	–	2006
<code>\sleep</code>	sleeping	Jan Wieck	2007
<code>\shell</code>	shelling...	Michaël Paquier	2009
<code>gaussian</code>	random	Mitsumasa Kondo	2014
<code>exponential</code>	random	Fabien Coelho	2014
<code>expression</code>	integer arithmetic	Robert Haas	2015
<code>double</code>	arithmetic and functions	Fabien Coelho	2016
<code>non-ascii</code>	variable names	–	2017

PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion

<code>-C</code>	connection	Tatsuo Ishii	2001
<code>-M</code>	query mode	Takahiro Itagaki	2008
<code>-T</code>	run time	–	2008
<code>-j</code>	threading	–	2009
<code>-R</code>	throttling	Fabien Coelho	2013
<code>-L</code>	latency limit	–	2014
<code>-f/b</code>	... weighted scripts	–	2016

PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion

<code>-l</code>	logging	Neil Conway	2002
<code>-r</code>	per statement stats	Florian Pflug	2010
<code>--sampling-rates</code>	sample stats	Tomas Vondra	2012
<code>--aggregate-interval</code>	aggregated stats	–	2013
<code>-P</code>	progress	Fabien Coelho	2013
<code>-f ...</code>	per script stats	–	2016

PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion



PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion

Initialize a database

`-i -s 1000 ...`

- create and fill, with scaling
- options: PK, FK, unlogged, fillfactor, tablespace...

Run scripts

`-T 1000 -c 32 -j 8 ...`

- psql-like, 3 builtins or custom, weighted, prepared, throttled
- parallelism: threads, clients, re-connections...

Measure and report performance


`-r -l -P 1 ...`

- tps, latency, timeout; per script, per command...
- detailed, sampled or aggregated; stdout or file

Beware

- long enough
- several times
- representative

*warm-up, checkpoint and vacuum
reproducibility
pedal-to-the-metal?*



Benchmarking vs Performance Testing

PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future

Needs
CommitFest
Conclusion

Benchmarking

System comparison

- standard schema and transaction
- maximum load
- report transaction per second
- latency should good enough. . .

pedal to the metal

tps

s

Performance Testing


Does it work for me?

- YOUR schema and transaction
- YOUR load. . .
- load must be processed
- latency must match application constraints

throttling

Version 9.5.5


■ throughput 329.4 tps
■ average latency 24.3 ms



■ latency stddev. 79.5 ms

Version 9.6.1

■ throughput 326.4 tps
■ average latency 24.4 ms



■ latency stddev. 20.3 ms

PgBench WIP

F. Coelho

PgBench

- History
- Capabilities
- Caveats


Performance

- Overheads
- Loading
- Connection
- SSL
- Index
- Fill Factor


Future

- Needs
- CommitFest
- Conclusion

Version 9.5.5



Version 9.6.1



What is happening?

- transaction surges are absorbed
- then data are written disk

Buy Now, Pay Later!

*in-memory + WAL
checkpoint*

PgBench WIP

F. Coelho

PgBench

- History
- Capabilities
- Caveats

Performance

- Overheads**
- Loading
- Connection
- SSL
- Index
- Fill Factor

Future

- Needs
- CommitFest
- Conclusion

Performance

PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future

Needs
CommitFest
Conclusion

Sleep zero

13.4 Mtps, 75 ns

```
\sleep 0
```

Set a variable

9.5 Mtps, 105 ns

```
\set i 0
```

Empty command

97,222 tps, 10.3 μ s

```
;
```

Empty SELECT

51,631 tps, 19.4 μ s

```
SELECT;
```


Impact of schema


on loading time

steps (d)rop (t)able (g)enerate (v)accum (p)rimary and (f)oreign key

■ pgbench -i -s 100 -I 'dtgv'	18 s
■ pgbench -i -s 100 -I 'dtgvp'	29 s
■ pgbench -i -s 100 -I 'dtgvpf'	32 s
■ pgbench -i -s 100 -I 'dtpgvf'	39 s
■ pgbench -i -s 100 -I 'dtpfgv'	103 s

Impact summary

■ Primary key	50-100%	■ Foreign key	20-300%
---------------	---------	---------------	---------



- Client 8 cores, 16 GB
- LAN 1 Gbps
- Server 16 cores, 32 GB, HDD

Initialization and Benchmarks

Postgres 9.6.1

<code>pgbench -i -s 100</code>	1.5 GB
<code>pgbench -T 2000 -C "host=server sslmode=require"</code>	36.1 tps
<code>pgbench -T 2000 -C "host=server sslmode=disable"</code>	56.4 tps
<code>pgbench -T 2000 "host=server sslmode=disable"</code>	105.4 tps
■ connection AAA	8.2 ms
■ SSL negotiation	10.0 ms
■ transfers and transactions	9.5 ms

SSL or not?

pgbench ... "sslmode=..."

PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future

Needs
CommitFest
Conclusion

SSL Costs

time & €

- negotiation and re-negotiation
- cryptographic functions
- certificate ?

Benefits

Snake Oil!

- Confidentiality
- Integrity
- Authentication

```
pgbench -j 1 -c 1 -D scale=100 -f ro3.sql -T 30 "host=server ..."
```

sslmode=require

SSL

- throughput *709.7 tps*
- latency *1.407 ± 0.132 ms*

sslmode=disable

clear

- throughput *781.6 tps*
- latency *1.277 ± 0.034 ms*

With primary key

17,225 tps

- initialization

```
pgbench -i -s 10 -I "dtgvp"
```

No primary key

23 tps

- initialization

```
pgbench -i -s 10 -I "dtgv"
```

With hash index

18,289 tps

- initialization

```
pgbench -i -s 10 -I "dtgv"
```

- plus non unique hash index

```
CREATE INDEX ah ON pgbench_accounts USING HASH(aid);
```

PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future

Needs
CommitFest
Conclusion

Update intensive load

with MVCC

- UPDATE = DELETE + INSERT
- induce about 3 page writes
- *or* keep some free space available

Initialization

```
pgbench -i -s 100 -F 95
```


```
CREATE TABLE pgbench_accounts(...)  
  WITH (FILLFACTOR = 95);  
...
```

Only UPDATE script

```

\set naccounts 100000 * :scale
\set aid random(1, :naccounts)
\set delta random(-5000, 5000)
UPDATE pgbench_accounts
  SET abalance = abalance + :delta, filler = NOW()::TEXT
  WHERE aid = :aid;

```



PgBench WIP

F. Coelho

PgBench

History

Capabilities

Caveats

Performance

Overheads

Loading

Connection

SSL

Index

Fill Factor

Future

Needs

CommitFest

Conclusion

Future

TPC-B Simple Banking Example

version 2.0, June 1994

- schema creation and initialization
- simple but non trivial SQL scripting
 - SQL commands (SELECT INSERT UPDATE COMMIT...)
 - variables, expressions, (uniform) random numbers, **if**...
- constant performance monitoring and reporting
stability, steady-state, statistics, errors...
- explicit anti-cheating constraints
*1.3.2 application **must retrieve** the balance*

Cannot be implemented with PgBench (yet)

Needed Capabilities

PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future

Needs
CommitFest
Conclusion

Missing features in development

<code>cset/gset</code>	get query results into variables	Fabien Coelho
<code>boolean</code>	comparisons and logical operators	—
<code>if/endif</code>	conditional expression	<i>refactoring done, waiting on boolean</i>

Other features in development

<code>pow</code>	another function	Raúl Marín
<code>ppoll</code>	for handling more clients	Doug Rady
<code>stats</code>	initialization step timing	—
<code>perm</code>	pseudo random permutation	...
<code>error</code>	handling...	?

TPC-B *Real* Transaction Profile

PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future


Needs
CommitFest
Conclusion

```
\set tbid random(1, :scale)
\set tid 10 * (:tbid - 1) + random(1, 10)
\if :scale = 1 OR random(0, 99) < 85 -- same branch
  \set bid :tbid
\else -- other branch
  \set bid 1 + (:tbid + random(1, :scale - 1)) % :scale
\endif
\set aid :bid * 100000 + random(1, 100000)
\set delta random(-999999, 999999)
BEGIN;
UPDATE pgbench_accounts
  SET abalance = abalance + :delta WHERE aid = :aid
  RETURNING abalance AS balance \gset
UPDATE pgbench_tellers
  SET tbalance = tbalance + :delta WHERE tid = :tid;
UPDATE pgbench_branches
  SET bbalance = bbalance + :delta WHERE bid = :bid;
INSERT INTO pgbench_history (tid, bid, aid, delta, mtime)
  VALUES (:tid, :bid, :aid, :delta, CURRENT_TIMESTAMP);
END;
```

Work in **slow** Progress

`commitfest.postgresql.org`

- submit patch
- **get a review** and review others
- get a decision
- **wait for a committer...**
- and maybe get one



PgBench WIP

F. Coelho

PgBench

History
Capabilities
Caveats

Performance

Overheads
Loading
Connection
SSL
Index
Fill Factor

Future

Needs
CommitFest
Conclusion

Easy to use and improving tool

pgbench

- write your custom script
- run it against your data
- for your load

Need test data?

datafiller

- directives on SQL declarations
- generators for many types and constraints