




2ndQuadrant +
Professional PostgreSQL

State of PostgreSQL Database 2015

Gianni Ciolli

5 4 3 2 ... MeetUs!
Milan, 12-13 May 2015




Notice

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 318633.


Seventh Framework Programme
FP7/2007-2013

Communications Networks, Content & technology



What is PostgreSQL?

- 19yo Database Management Server
- Relational
 - model founded on Set Theory
 - that's why it doesn't age! ; -)
- “Post-Relational”
 - Advanced features
 - Very Extensible
 - DBA-friendly & DEV-friendly



PostgreSQL Timeline

- Seasoned


- 9.3 9.2 9.1 9.0 ...
 Sept 2013 Sept 2012 Sept 2011 Sept 2010

- Stable

- 9.4
 - Sept 2014 Dec 2014

- Future

- 9.5
 - Sept 2015 (?)




Portability

- Architectures

As of today: x86, x86_64, IA64, PowerPC, PowerPC 64, S/390, S/390x, Sparc, Sparc 64, ARM, MIPS, MIPSEL, M68K, and PA-RISC.


- OS

Linux, BSD, UNIX, Windows, Mac...




Fully Open Source

- BSD licence
- Distributed Copyright
- No takeovers!
- Forever Free...




Transaction subsystem

- Fully ACID
- Durability ↔ Speed
 - Transaction-level
 - Zero Data Loss with Synchronous Replication
- Transactional DDL
 - No “deploy rollback” scripts needed!
 - Rehearse the deploy on live system...




Pluggable Languages

- Functions in different languages
 - C, SQL, PL/pgSQL, Perl, Python, PHP, Java, Javascript, Ruby, R...
- Creative use cases
 - PL/Proxy
Achieve large scalability via transparent sharding
 - PL/sh
Run your shell scripts from the database!



Pluggable Languages (cont.)


- Move code across client and server
 - You are free to choose the *best* architecture!
- No need for expensive translations...
 - You are free to choose the *right* language!



JSON, XML


- Partially- or Non-relational models
 - ACID without Sets
- Server-side parsing
 - Dedicated JSON and XML operators

```
SELECT xpath('/a/b/c/text()', xml, ...)  
       AS value_of_c  
FROM companies  
WHERE json->'hq' ->>'country' = 'UK';
```




Replication

- Streaming
 - Fully secure
 - Native OpenSSL
- Synchronous
 - Zero Data Loss
 - Transaction-level
 - Choose durability at commit!




Point-In-Time Recovery

- “Continuous Backup”
- Disaster Recovery
- Stop at a given timestamp or transaction Id
 - Maximum granularity
- Pause, resume
- Fix user mistakes




Foreign Tables

- Readable and Writable
- SQL-standard Foreign Data Wrappers
 - Pluggable adaptors
- Federate with other PostgreSQL servers
- Federate with other kinds of databases
- Federate with other data sources
 - Web services
 - Text files
 - ...




Features in 9.4

- Logical Decoding
 - Basis for future Multi-Master
 - Currently patch to 9.4
 - It's actually live... ???
 - we'll see later...
 - 9.4 supports Single-Master Replication
 - Flexible
 - Works with 9.4
 - Low overhead
 - Many use cases




Features in 9.4

- Unstructured data
 - Improved JSON
 - Improved GIN Indexes
 - Faster!
- Performance
 - Improved Materialised Views
 - Improved Write Concurrency
- Extensibility
 - Dynamic Background Workers




Work In Progress for 9.5

- BRIN Indexes
- Row-Level Security
- Read/Write Performance improvements
- SKIP LOCKED
- Parallel VACUUM
- INSERT ... ON CONFLICT DO NOTHING/UPDATE
 - A.k.a. “upsert”




BRIN Indexes

- Block Range INdices
- For very large tables
- “Automatic Partitioning”
- Very small index size
 - Fully in memory even with *huge* tables
 - Can be 10^{-5} of table size!
 - This is *a lot!*
 - Example: 50 TB table... 500 MB index




Future: BDR

- Bi-Directional Replication
 - PostgreSQL's Multi-Master
- In production!
- Composed of several patches
 - Some in 9.3
 - Some in 9.4
 - Some already in 9.5
 - Some hopefully in 9.5...
- For more information
Craig Ringer, blog.2ndquadrant.com, Dec 2014




Future: The AXLE Project

- axle-project.eu
- Advanced Analytics for Extremely Large European Databases
- Funded by EU 7th Framework Programme
- Partners
 - 2ndQuadrant
 - Barcelona Supercomputing Center
 - Portavita
 - The University of Ljubljana
 - The University of Manchester




AXLE achievements

- BRIN Indexes
- Statistics improvement
- Row Level Security
- Advanced Hardware Architectures
(FPGA Query Acceleration)
- ...




And now...

Questions?




And then...

Thank you!



Licence

This document is distributed under the **Creative Commons Attribution-Non commercial-ShareAlike 3.0 Unported** licence



A copy of the licence is available at the URL

<http://creativecommons.org/licenses/by-nc-sa/3.0/>

or you can write to

*Creative Commons, 171 Second Street, Suite 300,
San Francisco, California, 94105, USA.*