

2ndQuadrant

Professional PostgreSQL



Getting ready for PostgreSQL 9.1

Gabriele Bartolini

<Gabriele at 2ndQuadrant.com>

<http://www.2ndQuadrant.com/>



PostgreSQL


License

- Creative Commons:
 - Attribution-NonCommercial-ShareAlike 3.0
 - You are free:
 - to Share – to copy, distribute and transmit the work
 - to Remix – to adapt the work
 - Under the following conditions:
 - Attribution: You must attribute the work in the manner specified by the author or licensor
 - Non-Commercial: You may not use this work for commercial purposes
 - Share Alike: If you alter, transform, or build upon this work, you may distribute the resulting work only under the same or similar license to this one

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

About myself

- Founder of 2ndQuadrant Italia
- Database architect at 2ndQuadrant
 - Business Critical OLTP databases
 - Data warehousing
- PostgreSQL user for over 10 years
- Founder of PostgreSQL Europe
 - Former Vice President
- Founder of the Italian PostgreSQL Users Group
 - Current President



About 2ndQuadrant

- Professional Open-Source Software House
- Platinum sponsor of PostgreSQL
- 100+ customers worldwide
- Directly represented in 5 countries
 - Special coverage in Latin America and Nordic Countries
- 24x7 support for PostgreSQL production environments
- What we do
 - Design, Develop and Support PostgreSQL
 - We deliver high quality professional services
 - Leaders in HA, Recovery, Scalability and Performance
 - Training



Who's using PostgreSQL?




Table of contents

- The PostgreSQL Project
- PostgreSQL 9.0
- PostgreSQL 9.1




Part I

The PostgreSQL Project




Postgres or PostgreSQL?

- Officially it is **PostgreSQL**
- Postgres was the initial name of the project
- With SQL support (1996) the project changed name into PostgreSQL
- **Postgres** is commonly accepted as a synonym



PostgreSQL and its key features

- Open-Source Database Management System
- Close support of the SQL Standard (SQL:2008)
- Multi-Version Concurrency Control
 - Multi-user, designed for concurrency
- Extensibility:
 - Data types
 - Functions



Applications, databases and DBMS


Software applications

Software applications



100% Open-Source and Community

- Distributed under a BSD-like license
 - The “PostgreSQL License” (TPL)
 - <http://www.opensource.org/licenses/postgresql>
 - Highly **permissive**
 - Not subject to change of licensing
 - It will always be open-source
 - Not subject to acquisitions
- A community project



The “PostgreSQL” ecosystem

- Community
 - Developers
 - Users
- End users
 - Companies and organisations
 - Consulting companies
 - Individuals
- Sponsors
 - Development
- External Stakeholders
 - Companies that sponsor the development of key features



A mature project

- PostgreSQL has been developed for over 25 years
 - It is an adult
 - It is even allowed to drink (in some countries)
- It has grown rapidly in popularity
 - Particularly in the last 3 years
 - OpenStreetMap moved to Postgres (2009)
 - Provider of support for Solaris changed (2010)



PostgreSQL releases

- Current stable release: 9.0.4 (18 April 2011)
- Current Major release: 9.0 (September 2010)
 - 5 major releases in the last 5 years
- Oldest supported release: 8.1 (2005)
 - On Windows: 8.2
- Major releases supported by the community:
 - 8.1, 8.2, 8.3, 8.4 and 9.0
- PostgreSQL 9.1 is expected later this year
 - Currently 9.1.0beta2 is out




Part II

PostgreSQL 9.0



Architecture

- Client / Server
- Multi-processing
- ANSI C
- Uses stable system functions
- Multi-platform



Typical use cases

- Transactional databases (OLTP)
- Telecommunications
- E-Commerce
- Monitoring of electronic devices
- ERP, CRM, data warehousing, BI
- Social networking platforms
- High traffic websites
- Geographical Information Systems (PostGIS)
- Web CMS, Blog systems
- etc.




A versatile product

- Multiple interfaces for connectivity:
 - Native, JDBC, Python, Perl, PHP, ODBC, .Net, ...
- Multiple server languages:
PL/pgSQL, PL/Python, PL/Perl, PL/Java, ecc.
- User, path, options, environment



Cool robust features

- Views
- Stored procedures
- Triggers
 - Column and/or conditional
- ACID transactions
- Schemas
- Tablespaces
- Partitioning
- Custom types
- Rules



Security and data protection

- Concept of “security by default”
- Multi-level security model
 - server, database, table, column
- SSL
- Integration with the company's security infrastructure
- Full disclosure of vulnerabilities
- Rapid fix of security bugs



Features for web technologies

- Support for SQL:2008
- Support for the UNICODE standard (UTF-8 charset)
- Support for regular expressions
- Support for XML
- Database level collation
- Localisation



Advanced web features

- Full text search
- Arrays
- Transparent compression of text data
- Data types (enums, uuid, cidr, macaddr, hstore, ...)
- Session parameters with SET
- Asynchronous commit



Business critical environments

- Two of the major challenges and requirements:
 - Maximise the uptime (“5 Nines” = ~ 5 minutes/year)
 - Reduce downtimes
 - Remove data loss
 - While keep data consistent (always)



PostgreSQL's answer

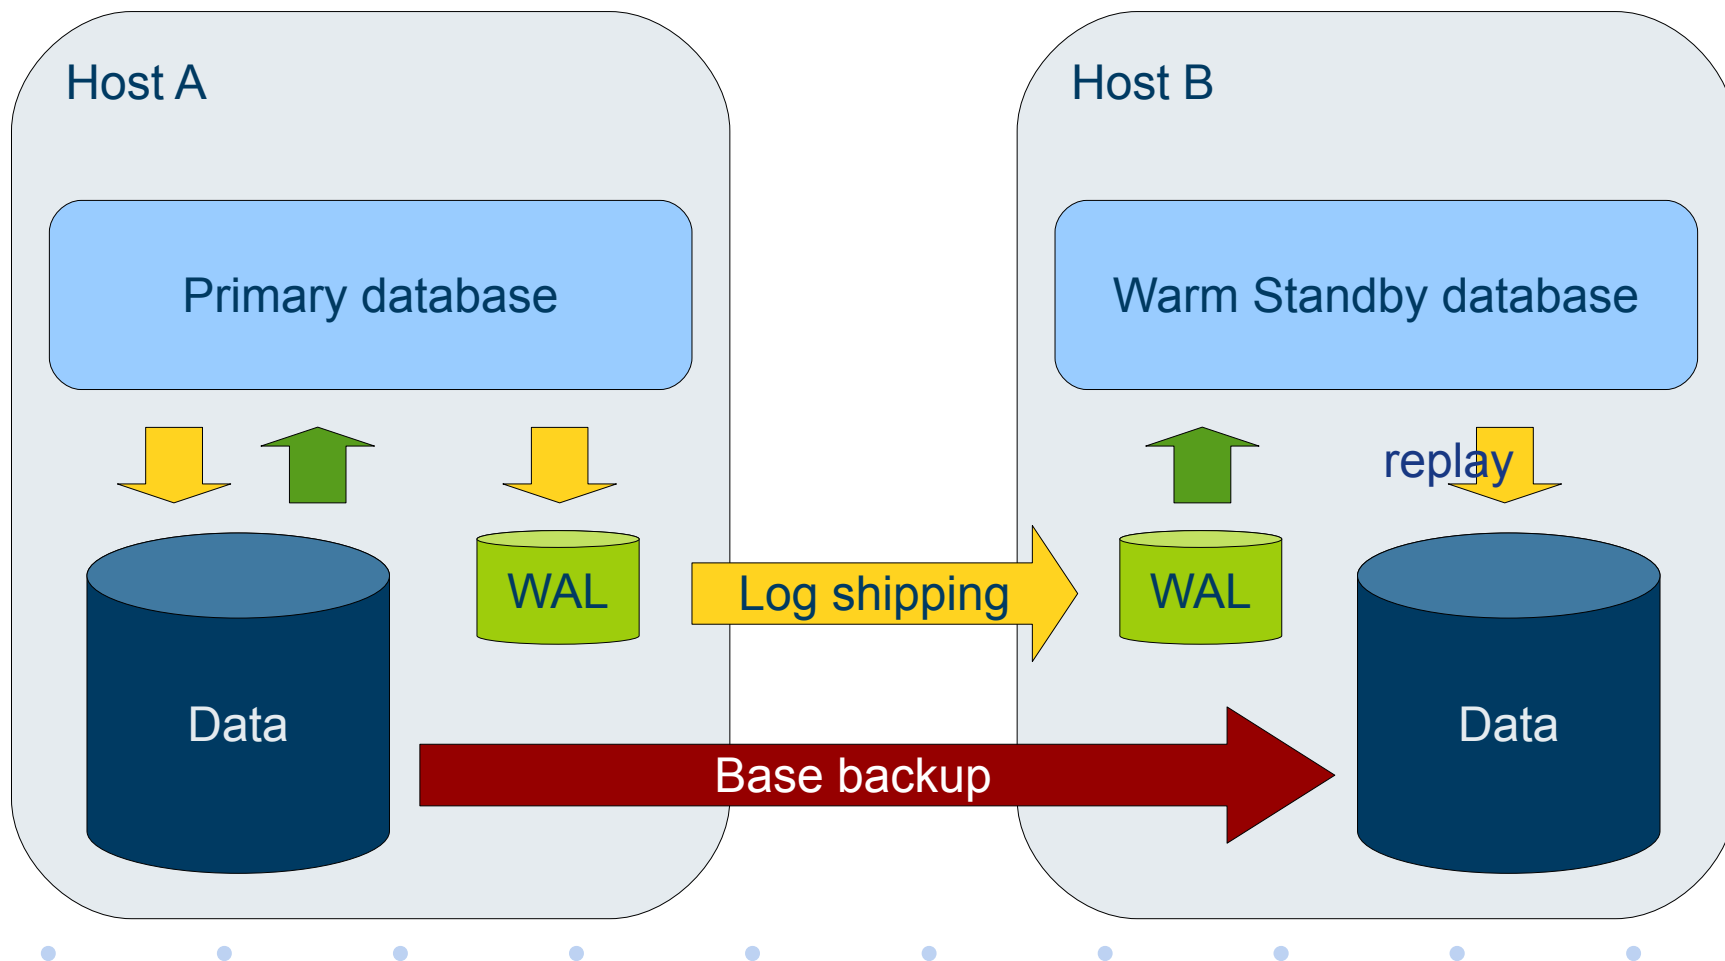
- Robustness and reliability
 - WAL Files
- Hot backup
- Online Backup and Point in time recovery
- High Availability
 - Warm/Hot Standby
- Asynchronous Streaming Replication
- Synchronous Streaming Replication (9.1!)




Warm standby

- Technology for **asynchronous replication** based on:
 - Write Ahead Log (WAL) files shipping
 - Through operating system commands (e.g. rsync)
- 1 primary node, N secondary nodes
- Secondary nodes:
 - Are called **standby** nodes
 - Do not accept connections
 - Are an almost identical copy of the master
 - Small delay due to WAL shipping
- High Availability technology
- The standby is promoted in case of **failover** (unexpected) or **switchover** (deliberate)

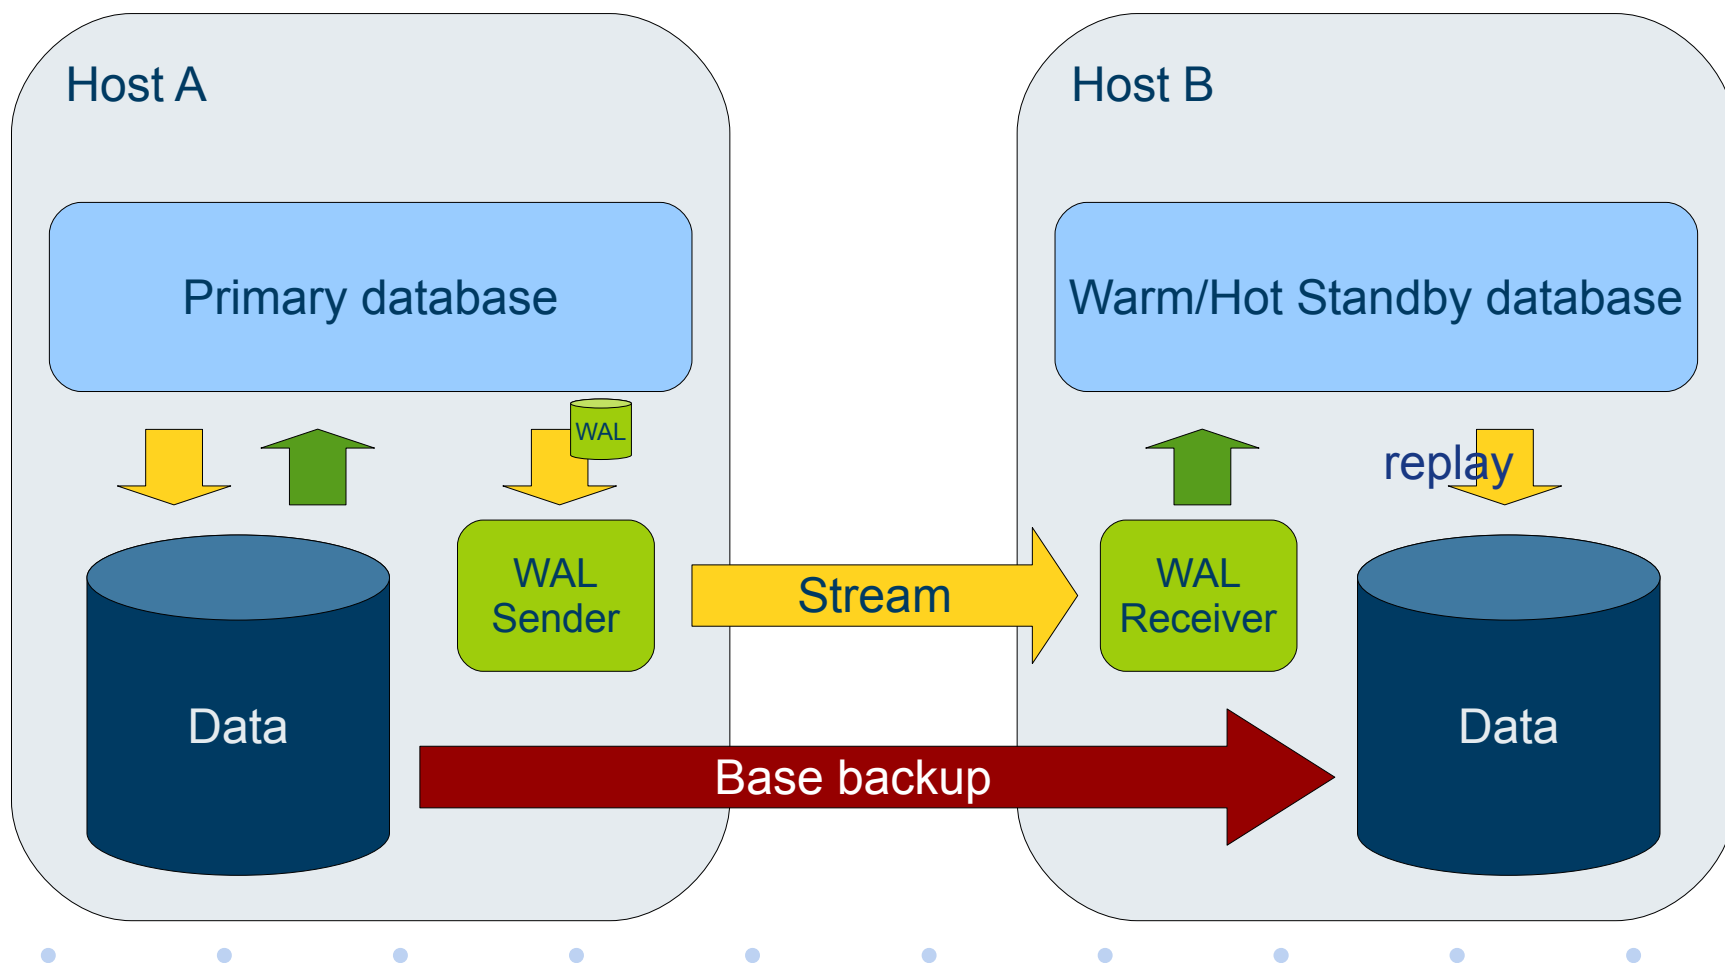
Warm Standby, overview



Hot Standby



Streaming Replication




More replication solutions

- Not part of PostgreSQL's core
 - Only available solutions for PostgreSQL < 9
- A few alternatives
 - Slony
 - PGPool
 - Bucardo
 - Londiste (Skype)




VLDB

- Many 1TB+ databases with PostgreSQL
- Typical usage is 100GB+ databases
- Key features:
 - SMP architectures
 - Tablespace
 - Horizontal partitioning



Scalability

- Caching:
 - pgmemcache - integration with memcache
- Distribution on multiple nodes:
 - PgBouncer (connection pooling)
 - PL/Proxy
 - Map/Reduce (RUN ON ALL, ANY, single node)




Administration and monitoring tools

- Command line:
 - psql, pg_dump, pg_restore, ecc.
- GUI:
 - PgAdmin3
- Web:
 - PHPpgAdmin
- Alerting:
 - Nagios integration
- Monitoring:
 - Cacti and Munin integration
- Replication:
 - repmgr (*devel*)



Spatial extension: PostGIS

- PostgreSQL's most important extension
- Geographical Information System (GIS)
- OpenStreetMap
- Add geographical data types
 - Open Geospatial Consortium
- Geodatabase




Part III

PostgreSQL 9.1



Overview

- Improvements in operations
- Extensions packaging
- New features and performance improvements in key areas
- **Synchronous Replication**
- Triggers on views (“updatable views”)
- Column collations
- Writable CTE
- Serializable Snapshot Isolation
- Better partitioning
- ...



Synchronous Replication

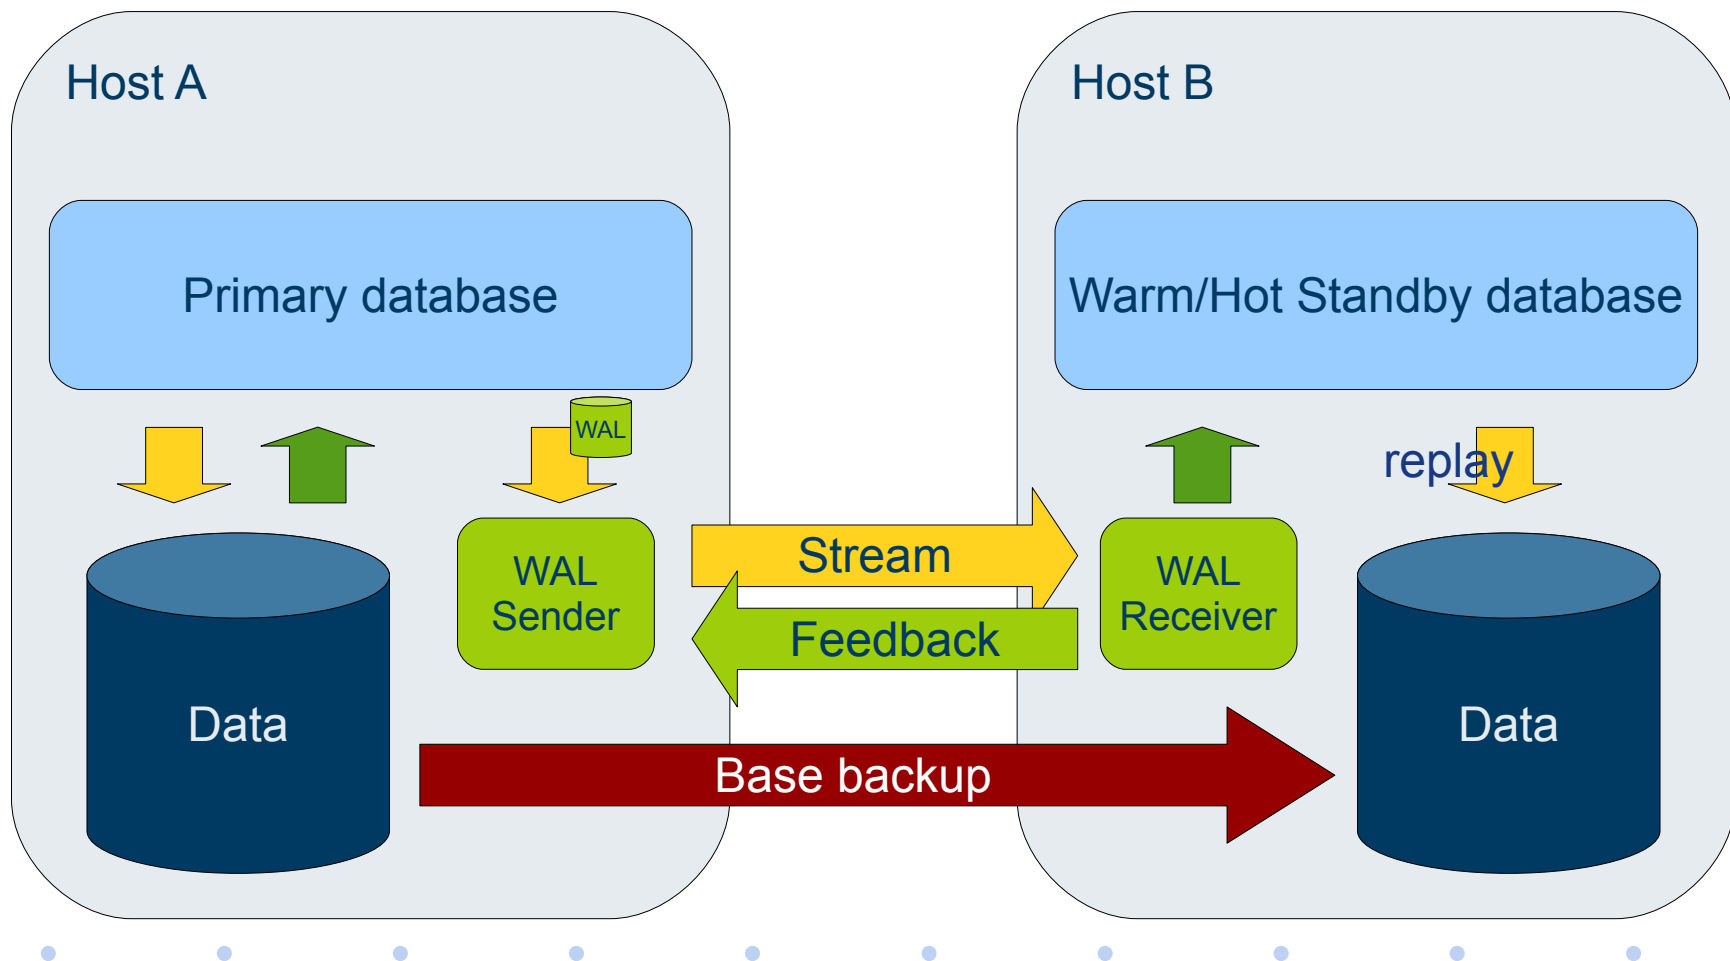
- Data loss prevention technology
- Uses streaming replication
- Performance downgrade
- Main concept:
 - A transaction is written on ≥ 1 standby servers before COMMIT
- Priority ordered list of servers
 - Identified by their `application_name`
 - `synchronous_standby_names = '*'`
 - `synchronous_standby_names = 'slave1, slave2, slave3'`
- Can be controlled at several levels:
 - Transaction, session, user, database, server
 - `SET synchronous_commit TO off;`

Example of transaction control

- **SET synchronous_commit TO off;**
- BEGIN; *-- I could not care less*
- INSERT INTO bands ('Take that');
- INSERT INTO bands ('Black Eyed Peas');
- COMMIT;
- **SET synchronous_commit TO on;**
- BEGIN; *-- Extremely important!*
- INSERT INTO bands ('AC/DC');
- INSERT INTO bands ('Cream');
- COMMIT;




Synchronous Replication



A few improvements from 9.0

- pg_basebackup
- Improved monitoring
 - Views, including `pg_stat_replication`
- Feedback from Hot Standby servers to the master about running read-only queries
 - Prevents query cancellation (VACUUM on master)
- Allows to pause/resume replication on slaves




Extensions before 9.1

- Extensibility has been a key feature in PostgreSQL
- With one limitation:
 - No dump/restore support
 - No special attention for their objects
 - Issues with upgrades
- `pg_dump` cannot distinguish extensions' object from the rest of the objects
 - Extensions, including functions, are included in the dump
 - Sometimes this can be a mess



Extensions from now on

- `pg_dump` is now able to produce only this:
 - `CREATE EXTENSION IF NOT EXISTS extname`
- Programmers can define their own extensions through a control file, which specifies metadata and **contained objects**
 - Version, SQL files directory, modules path, dependencies, ...
- **More commands:**
 - `DROP EXTENSION extname`
 - `ALTER EXTENSION extname UPDATE`
- **PGXN: PostgreSQL Extension Network (pgxn.org)**



Example of extensions: PGMP

Homepage: <http://pgmp.projects.postgresql.org/>

pgmp.control file

```
default_version = '1.0'  
comment = 'Multiple Precision Arithmetic extension'  
directory = 'pgmp'  
relocatable = cd true
```



Column level collations

- Collations control how to perform character comparisons within strings
- Crucial in non English speaking countries
 - Think about Europe!
- You can now specify different collations inside the same table, at **column level**

```
CREATE TABLE collate_example (  
    c_text TEXT COLLATE ucs_basic,  
    en_text TEXT COLLATE "en_EN.utf8",  
    it_text TEXT COLLATE "it_IT.utf8"  
);
```



Example of collations

- `SELECT it_text FROM collate_example ORDER BY c_text;`

a
e
i
à
é
...

- `SELECT it_text FROM collate_example ORDER BY it_text;`

a
à
e
é
i
...



Triggers on views

- `CREATE TRIGGER trigger_name INSTEAD OF event ON view_name FOR EACH ROW;`
- Events:
 - INSERT
 - UPDATE
 - DELETE
- Allows views to be **UPDATEABLE**



Unlogged tables

- CREATE UNLOGGED TABLE
- Do not produce any WAL content
 - Not crash safe (`TRUNCATE` on recovery)
 - Their content is not replicated
 - Just the schema
- Typical use case:
 - **Caching**
 - Storage of Web temporary sessions
- Differences with temporary tables
 - They go beyond a session's lifetime
 - Global namespace and visibility



Foreign tables

- Part of SQL/MED
 - Prior to 9.1: dblink, DBI-link, PL/Proxy
- Foreign data wrapper API (FDW)
- CREATE FOREIGN TABLE
- Currently only one example as contrib
 - file_dw (foreign data wrapper for reading CSV files)
 - Exposes a CSV as a FOREIGN TABLE in the database
 - Supports only SELECT
 - Mainly a test case to show the potential of the API
- It will be improved from 9.2



Writeable Common Table Expressions

- Improvement to WITH queries introduced in 8.4
- Support for write queries within the WITH clause

```
WITH d AS (  
    DELETE * FROM log WHERE logday = '2011-06-01'  
    RETURNING *  
)  
INSERT INTO old_log SELECT * FROM d;
```



Other important changes

- Serialisable Isolation Level
 - Great feature (too complex to talk about it now!)
- Optimization of queries that use ORDER BY, LIMIT, or MIN/MAX with inherited tables
- K-Nearest Neighbour GIST
- Standard conforming strings
 - By default: \ (escaping) throws an error (use ")
- `format()` (for `sprintf` lovers!)
- SE-Linux Integration for PostgreSQL
- ... and more
 - For a full list of changes, see the Release Notes


Conclusions

- PostgreSQL allows a database to be:
 - Fast
 - Consistent
 - Reliable
 - Secure
 - Integrable
 - Available
 - Recoverable
- ...



... and reduces the costs

- No license fee
- Simple license model
 - No dual licensing headaches
- Reduces the total cost of ownership (TCO) for a database solution
- Invest in knowledge
- Removes any vendor lock-in / monopolisation
- Largest community for an open-source database management system project



Let's make the Elephant happy!



Questions?



Thank you

- You can contact Gabriele via email at Gabriele@2ndQuadrant.com
- For more information on our professional services on PostgreSQL and Greenplum, visit our website <http://www.2ndQuadrant.com/>
- See you next year!

