PostgreSQL 9.0 Changes 4/11/2011

PostgreSQL version 8.5 has been renamed version 9.0. A major feature of 9.0 will be the addition of built in log-streaming replication. This allows logs to be streamed to the slave over the network which greatly increases the frequency at which these logs can be sent. Version 9.0.0 was released in September 2010.

psql
- improved how psql displays wrapped-around records
- prevent overwriting of command history if two psql sessions are running

vacuum
- improved "vacuum –full"
  - faster
  - automatically reindexes
  - solves problem of "index bloat"

new functions
- pg_table_size and pg_index_size to make gathering size info easier

triggers
- added column trigger (trigger executes when column is updated) and when trigger (trigger executes when simple IF-THEN conditions are met)

replication
- now allows write-ahead log (WAL) files to be streamed across a network connection to a standby server

PL/Python
- added support for Python 3

PL/pgSQL
- language installed by default
- no longer allows certain variable names which match certain SQL commands
  - these variable need to be double quoted

migration tool
- new pg_upgrade utility
  - utility for upgrading a database from 8.3 - 8.4 or from 8.4 - 9.0 in-place
  - replaces the dump/restore operation
    - dump/restore can still be done but is much slower
  - previously called pg_migrator
  - major bug fix in version 9.0.4
  - cannot be used with versions older than 8.3 !!!

\copy
- allows parameters to be specified within parentheses

- old syntax still allowed

SQL

- use of “=>” operator is deprecated
    - will be removed in future versions
- changed to allow statements with “IS NOT NULL” to use indexes
    - useful in statements which use MAX() or MIN() on columns containing NULLs

DDL

- new GRANT/REVOKE IN SCHEMA makes mass permission changes easier

ecpg

- added support for SQLDA (Descriptor Area)
- allow use of “new’ and “old” variable names

indexes

- can build indexes that include ordering (e.g. desc NULLS LAST)

MAX/MIN

- fixed problem where use of MAX and MIN in a partitioned table resulted in a sequential scan

Server Configuration

“Unfortunately version 9.0 only made server configuration more complicated. And newer Linux kernels even pushed default behavior backwards. Starting with Linux kernel 2.6.33, the default value picked for wal_sync_method changed to open_datasync. This turns out to have terrible performance implications for PostgreSQL, particularly when combined with the low default setting for wal_buffers in the server.”