

# Pgpool-II Development Status Updates

Pgpool-II Global Development Group

# We are moved!

- Due to closing of pgfoundry, we have moved to new web site(November 2011)
  - <http://www.pgpool.net>
    - Hosted by SRA OSS, Inc. Japan
  - Provided services
    - Mailing lists
    - Source code downloading
    - Documentations and Wiki
    - Source code repositories are now hosted by [git.postgresql.org](http://git.postgresql.org)
    - Many thanks to Guillaume Lelarge and Magnus Hagander to help migrating from CVS to git

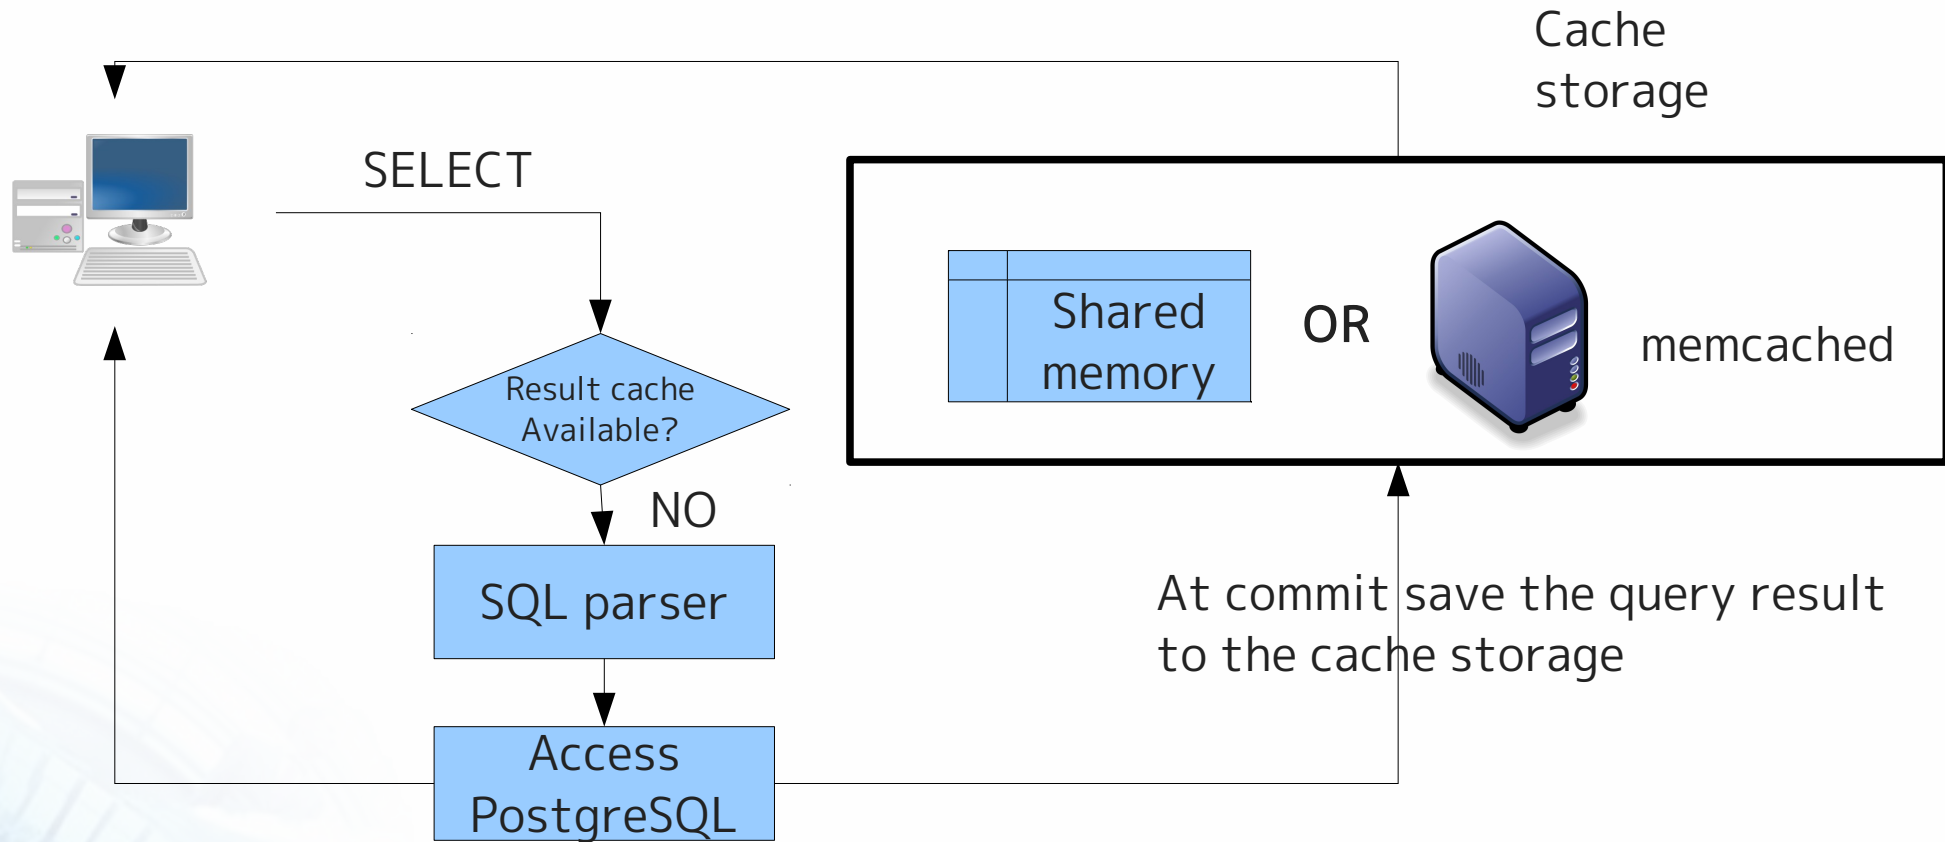
# Releases

- Stable releases
  - Pgpool-II 3.1 series
    - The first release was Pgpool-II 3.1.0(2011/9/8)
    - The latest release at this point is Pgpool-II 3.1.3 (2012/4/23)
    - About 5,000 downloads so far
  - Pgpool-II 3.0 series
    - The first release was Pgpool-II 3.0.0(2010/9/10)
    - The latest release at this point is Pgpool-II 3.0.7(2012/4/23)

# Current development

- Developing pgpool-II 3.2
- Major feature is “on memory query cache”
  - Similar to MySQL's query cache
  - Implemented in pgpool-II, rather than in database engine has some benefits
    - No access to database engine at all if cache hits
    - Can work with wide variety of PostgreSQL version (even possible with ancient 7.4)

If there's cache for the query result,  
returns it directly to client



If there's no cache for the  
query result, get result from  
PostgreSQL and returns to client

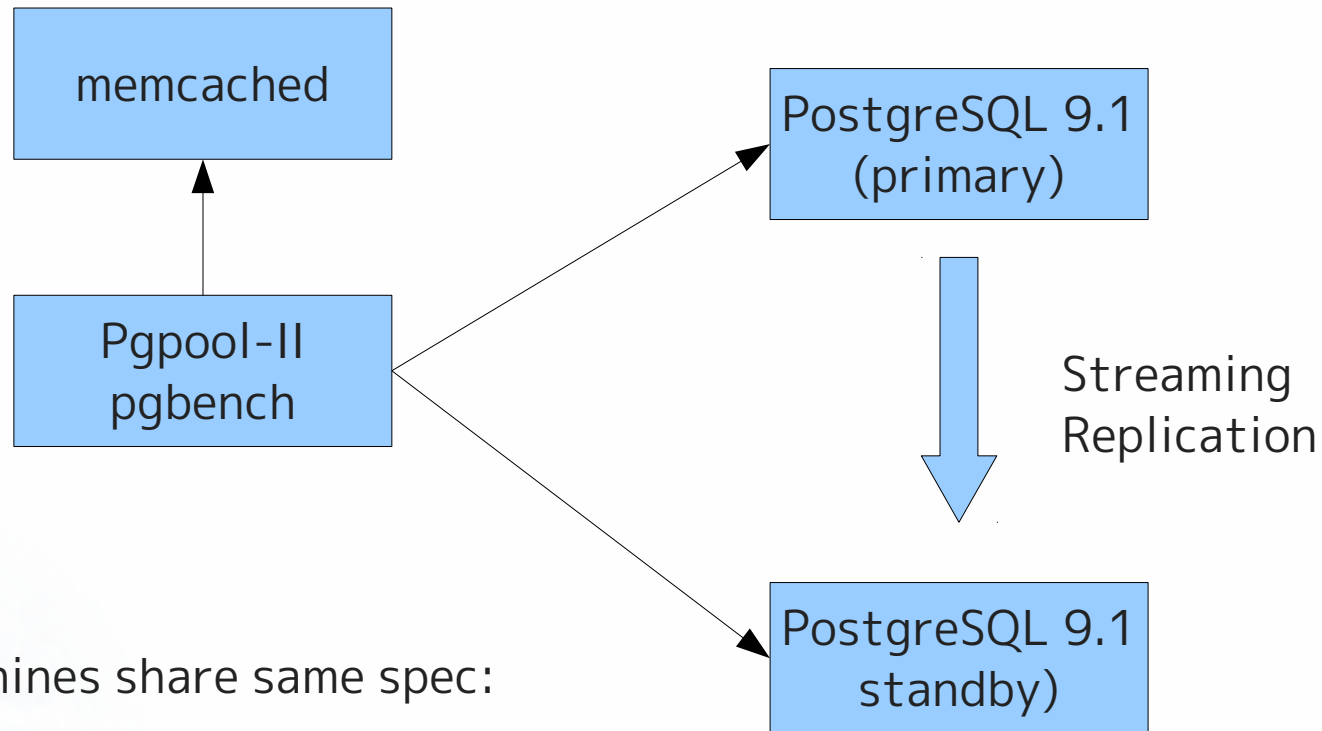
# Major features of on memory query cache

- If cache hits, no access to PostgreSQL is involved, thus very fast
- Choices of cache storage
  - Shared memory
    - Fast and easy to use but cache size is limited
  - Memcached
    - Can be used as very large cache storage, but network overhead is involved
- Automatic cache invalidation
  - Table updates (including ALTER TABLE)
  - Cache expiration by specified duration
- Can specify which table is cached and which is not

# Limitations

- If a table is modified, all caches for the table are invalidated
  - Not recommended for update intensive systems
- Users can only access to cache which was created by the same user(security reason)
- Views are not cached
  - Pgpool-II is unable to know the updates to the underlying tables
- Cascade update to tables are not recognized
- Table update by trigger is not recognized
- Too large query result cannot be cached(~1MB, depending on the configuration)

# Benchmark configuration




All the machines share same spec:  
Cent OS 5.5  
Core 2 Duo 2.33GHz  
Mem: 2GB



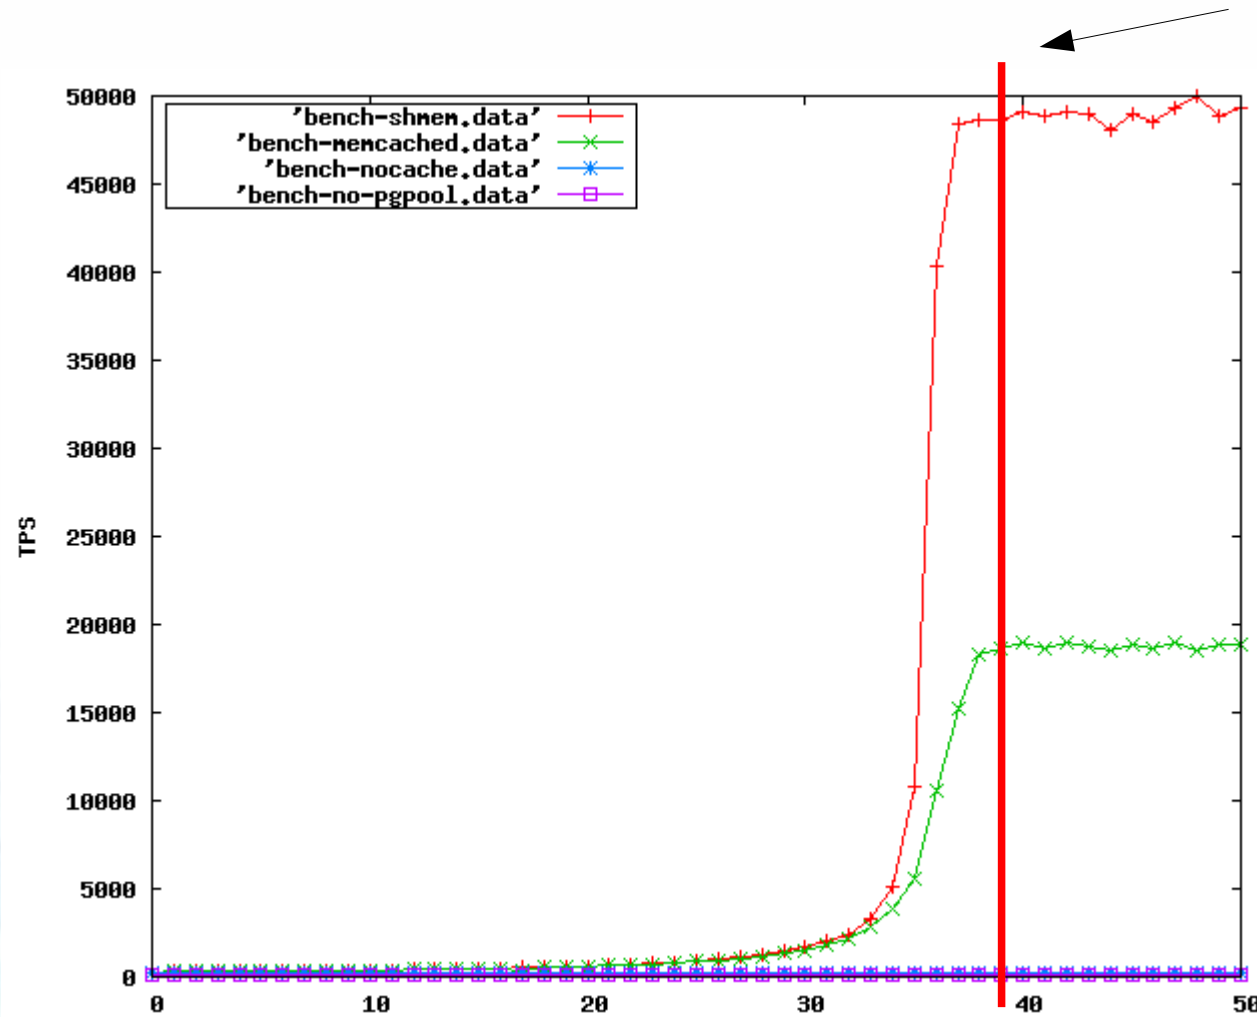
# Benchmark configuration

- Scale factor = 1(13MB DB)
- pgbench\_accounts primary key is removed to simulate non trivial SELECTs
  - Each SELECT takes about 6.7ms
- SELECT only test
  - Pgbench -S -T 10 (1 run is 10seconds)
- Cache storage is 64MB(for both shmem and memcached)

# First 20 run of pgbench



# Total run of 50



- After cache is filled up both cache method keep high performance.
- Shmem outperforms memcached significantly.

# Milestone/Schedule

- Pgpool-II 3.2 alpha1 expected to be released this week
- Pgpool-II 3.2 beta1 is expected to be released by the end of May
- Pgpool-II 3.2 official release is expected in June
- Tests and documentations(French and Chinese) will be the major work before the 3.2 release!