



PostgreSQL 8.3 Performance Features

Simon Riggs
2nd Quadrant
simon@2ndQuadrant.com




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




Performance Features

- Enhanced Scalability
- HOT
- Database Size Reductions
- Smoothing
 - Checkpoint smoothing
 - Xlog File Switch tuning
- Performance Logging
- Async Commit
- COPY tuning
- Recovery I/O Reduction
- Sort Optimizations
- L2 Cache tuning
- Synchronous Scan




MySQL v PostgreSQL Scalability

Tweakers.net Database-simulatie
MySQL 5.0.20a vs. 5.0.32-bk




- 5.0.20a (single)
- 5.0.20a (1x dual)
- 5.0.20a (1x quad)
- 5.0.20a (2x quad)
- 5.0.32-bk (single)
- 5.0.32-bk (1x dual)
- 5.0.32-bk (1x quad)
- 5.0.32-bk (2x quad)

Tweakers.net Database-simulatie
Clovertown vs. Woodcrest - PostgreSQL 8.2-dev




- Woodcrest 3,0GHz (single)
- Woodcrest 3,0GHz (1x dual)
- Woodcrest 3,0GHz (2x dual)
- Clovertown 2,66GHz (single)
- Clovertown 2,66GHz (1x dual)
- Clovertown 2,66GHz (2x dual)
- Clovertown 2,66GHz (1x quad)
- Clovertown 2,66GHz (2x quad)





MySQL v PostgreSQL Scalability


<http://jamonation.com/node/734>

pgsql sysbench on 8-core opteron, FreeBSD 7.0 with N CPUs active




mysql sysbench on 8-core opteron






Heap Only Tuples (HOT)

- Frequent updates cause mini-VACUUMs to occur on blocks as-needed
 - No need to wait for a full table VACUUM
 - Provides consistent performance
- Mechanism: Don't store index entries when we update a row that fits on the same block
- Only if we don't update an indexed columns
- 18 month R&D project, 8 people, ~4 man years




Async Commit

- Option to not wait for disk write when we COMMIT
- Can provide as much as 900% performance gain
- Can be applied on selected transactions only
 - Can mix async commit and standard synch-commit
- More flexible version of MySQL “speed-rather-than-robustness” feature




Scan Performance

- L2 Buffer Cache Recycling
 - Reuse small circle of buffers for SeqScan & VACUUM
 - Improves L2 cache hit ratio: +20% performance
- Synchronous Scans
 - Multiple sessions re-use the same buffers while scanning, providing a huge I/O reduction



PostgreSQL 8.3

- Smaller
- Faster
- Smoother
- More scalable
- More tunable



PostgreSQL