

Postgresql Performance

Compared to Informix (on Linux)
At CBRFC

Methods

- I made a reasonable effort to level the playing field, but in the end it was not perfectly level and differences are described in the next two slides about configuration and conditions

Configurations


- Informix
 - Version 7.31 (Linux)
 - Machine is a dual Xeon 2.2Ghz w/1.7G ram
 - Modified configuration
- Postgresql
 - Version 7.4.6 built from source code (not RPM)
 - Machine is a dual Xeon 2.8Ghz w/6G ram
 - Modified configuration

Conditions of the test

- Informix machine was operating with normal operational load
- Postgresql machine was otherwise idle
- So the postgresql machine has the advantage here, being faster with more ram and less load


Performance – mass loading

- 1 million rows into one table
 - Informix, dbload, 1000 rows per transaction
 - Postgresql, psql, \copy




Performance - queries

- Time series query from a 5 million row table with no index – ID and data-type specified without date information



Performance – storage efficiency

- Informix – cooked file space – 37G ascii data took 69G of disk space
- Postgresql – 4.8G ascii data took 18G of disk space



Conclusions

- Mass loading performance is on par with Informix
- Storage efficiency loss can be solved with more disk space (may need 2x the disk space to store the same amount of data)
- Query performance may be similar to Informix (was ~equal when limit 1 was applied to select statement)

Conclusions, cont.

- Better Performance? That depends on where you are coming from
 - Moving from HP-Informix it will be an improvement
 - From Linux-Informix it may be a step back
- Better overall?
 - Open source, free
 - Arrays!

Conclusions, cont.

- Did I mention Arrays? !!!
 - A better, less awkward way (than CBRFC-fastetc method) to store time series data
 - May increase performance
- Maintenance & Fragmentation (extents)
 - Minimal!
 - No inter-table fragmentation with pg
 - Not more than one table per file-system file
 - Vacuum commands will have to be ran routinely. This is essential.