

RE: Tuning for PostgreSQL Database

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2003-07/0091.html>

From: Rod Taylor (rbt_at_rbt.ca)

Date: 07/26/03

To: Tom Samplonius <tom@sdf.com>

Date: 26 Jul 2003 05:32:57 -0400

> > *But now that I have looked a bit more I see that this link*
> >
> > http://www.postgresql.org/docs/aw_pgsql_book/hw_performance/node6.html
> > ...
> >
> > *I would go with the information from the PostgreSQL site. It is at*
> > *least reviewed by the developers (if not written by them).*

In this case Bruce is a developer, but the advice in his book is rather old (much has changed in PostgreSQL over the last couple of years).

Short answer is that you need to benchmark your specific situation.

The more writes you have, the lower you will want the shared buffers (PostgreSQL isn't particularly efficient with buffers in this manner). Between 4000 and 8000 is usually a good number for a busy primarily write based db.

In regards to read performance, higher numbers can help but each step up will help less than the previous. If all of a sudden things get busy and some other elements (sorts, aggregate hash, ...) something is going to swap if you have the buffer cache set too large. Once the system starts swapping it will be difficult for it to catch back up again.

75% of RAM is generally too high. It leaves little space for several large queries to do their work in RAM (remember, buffer cache won't automatically resize itself when space becomes short). For a system running one or two report style (large amounts of data involved, aggregates all over the place, etc.) queries it may be ideal.

Most suggestions put buffer cache around 25% of the ram size. It's a safe number that comes close enough to peaking in read performance without too much of a hit to write performance. Tom Lane still gives the advice that you not go over a 10k due to the write performance hit and the lack of solid proof that read performance is improved by larger values.

freebsd-performance: RE: Tuning for PostGreSQL Database

All of the above assumes a recent version of PostgreSQL and a modern operating system. 7.4 will probably (not yet benchmarked) require a slightly `buffer_cache` setting than 7.3 due to increased ram usage for the work involved (speed increase) and indexes do not grow unbounded, as such take less space.

Please join `pgsql-performance@postgresql.org` to continue the debate. Nearly every PostgreSQL DBA and developer has a different opinion about exact values, but they're usually within the same ballpark.

Of course, the real fun is the the sweet spot changes with your data (distribution and quantity), query load, hardware and operating system. Even if you find the best value for todays load, you will quickly find it no longer to be true -- which is another reason to err on the conservative side.

-
- application/pgp-signature attachment: This is a digitally signed message part