Re: Tuning for PostGreSQL Database

Source: http://unix.derkeiler.com/Mailing−Lists/FreeBSD/performance/2003−07/0085.html

---

From: Christopher Weimann (*cweimann_at_k12hq.com*)
Date: 07/25/03

```
Date: Fri, 25 Jul 2003 12:26:01 −0400
To: Tom Samplonius <tom@sdf.com>
```

On Fri 07/25/2003−12:03:32AM −0700, Tom Samplonius wrote:
>
> *Maybe you should continue to worry. PostgreSQL isn't MySQL (or a*
> *typical server application). It reads all database pages into its shared*
> *memory area. It is wasteful for the DBMS and the OS to both cache this*
> *data. You'll want the PostgreSQL shared memory size to be around 75% the*
> *size of RAM (on a dedicated DBMS server). In fact, many commercial DBMS*
> *systems will use raw writes to bypass the OS cache!*
>

I was concerned about the disk cache because of this link

http://www.varlena.com/varlena/GeneralBits/Tidbits/perf.html#shbuf

which says...

> PostgreSQL counts a lot on the OS to cache data files and
> hence does not bother with duplicating its file caching effort.
> The shared buffers parameter assumes that OS is going to cache
> a lot of files and hence it is generally very low compared
> with system RAM. Even for a dataset in excess of 20GB, a
> setting of 128MB may be too much, if you have only 1GB RAM
> and an aggressive−at−caching OS like Linux.

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

which says...

> Ideally, the POSTGRESQL shared buffer cache will be:
>
> * Large enough to hold most commonly−accessed tables
> * Small enough to avoid swap pagein activity

So I have conflicting documentation.

I have machine with 4Gig of ram. What is the maximum
value of SHMMAX on FreeBSD?