

## RE: FreeBSD 5.3 MySQL Performance

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2005-02/1519.html>

---

**From:** Ted Mittelstaedt (*tedm\_at\_toybox.placo.com*)

**Date:** 02/10/05

To: "Joe Kelsey" <joe@zircon.seattle.wa.us>, "Matt Olander" <matt@offmyserver.com>

Date: Wed, 9 Feb 2005 23:20:42 -0800

owner-freebsd-questions@freebsd.org wrote:

> *On Wed, 2005-02-09 at 16:44 -0800, Matt Olander wrote:*

>> *On Wed, Feb 09, 2005 at 07:58:56PM -0500, Alec Berryman wrote:*

>>>> *Also, does anybody have any FreeBSD 5.3/MySQL benchmarks? I*

>>>> *searched the mailing lists but didn't turn up anything.*

>>>

>>> *There was an article posted to Newsforge today about benchmarking*

>>> *MySQL on different operating systems.*

>>

>> *oh! thanks...but according to this article, Linux outperformed*

>> *FreeBSD in every metric shown :-(*

>>

No, not true, reread the article. Performance was on par for uniprocessor

versions of Linux and FreeBSD for most tests. Also, the author said the following about the testing:

"highest performer in one category for a limited set of tests does not a "best" operating system make."

Please keep in mind that editors of publications love benchmarking articles, because they always get someone's tit in a wringer, and attract a lot of attention. And attention sells newspapers.

But you shouldn't take these things too seriously. The article points out a few things and the accompanying reader responses point out a few more things that are educational if you are choosing to run a database on a UNIX system, but by no means should the article be used as the sole basis for choosing one OS over another.

People that do benchmarking and publish the results are generally hoping to help point out problems. Sometimes this is because they have an axe to grind and want to see their favorite OS or program or whatever get some attention, sometimes just because it's nice to see some of your work in print. But regardless of why they do it,

## freebsd-questions: RE: FreeBSD 5.3 MySQL Performance

the results are valuable, because if problems that benchmarking reveals weren't pointed out, they wouldn't ever get fixed.

My take on the article is the most surprising thing in it was that Sun's own support staff couldn't answer the authors query about why Solaris was so slow, and the author finally figured it out by himself (The filesystem wasn't mounted with the forcedirect option) and set the needed option, whereupon performance dramatically improved. We always hear from commercial OS vendors how their products are so much better because they are supported – well it seems to me that if Sun's support was this bad for their own OS, well that throws the entire argument out the window, don't it?

>> *is that accurate?*

>

> *LinuxThreads is the WORST implementation of threading that anyone can imagine. Do not ever use Linux or the horrid LinuxThreads for anything that you want to save.*

>

> *Any so-called "benchmark" comparing Linux to anything else (especially windoze) has been polluted by the tradition in the linux/windoze world of running their disks in the completely unsafe "asynchronous" mode so popular with the ATA disk drive manufacturers.*

The author of the article avoided this by using a test method that in his words:

"I performed one test run to prime the system, almost all of the data was cached by MySQL, so there was little or no disk access."

>

> *Many companies have used FreeBSD and MySQL for years and years. There is no reason to not jump to FreeBSD and start using MySQL.*

Exactly, we use MySQL and FreeBSD quite a lot and have no problem with it.

> *At my last*

> *job, we ran very large MySQL databases on FreeBSD. For speed we used 15,000 RPM SCSI-3 disk drives. This gives you all the speed you need with the guaranteed safety of FreeBSD. Of course, SCSI-3 15,000 RPM drives are more expensive than those wimpy ATA drives.*

>

You see, this here points out the crux of the problem.

Boiled down the article essentially said that Linux performed better because it's SMP implementation allowed mysql to take advantage of both CPU's while FreeBSD's SMP implementation didn't.

But you see the problem with this is that in a real life situation, it is not often that you have such a small database and such a large amount

of system memory that the OS can load the entire database into a disk cache in ram. As you can no doubt understand, if the database is on disk all the additional CPU's in the world won't make the database run any faster once the disk channel gets saturated, which is easy to do.

And even if you can load the entire database in ram, if you make a lot of writes to it, the system has to push these to the disk channel eventually, unless of course you like for your entire database to vanish if there's a power interruption or system crash of some kind. So for a situation of a steady stream of writes, you end up I/O bound again. SMP on a database is no help if the system is I/O bound.

And if your database is going to I/O bind, because of how it's used and setup and how big it is, then this benchmark article is completely useless to you.

And even if your database isn't going to I/O bind then read the following comment one of the readers posted regarding OpenBSD and FreeBSD:

"They both use userland only threading, and therefore mysql is only ever running on a single CPU, no matter how many are in the system ...Different processes will run on different CPUs just fine, postgresql for instance would see an increase in performance on these platforms"

One last thing about SMP – rarely is mysql used in a vacuum, many times it's paired up with apache. In those cases if there's multiple accesses to the server at the same time then even though mysql only runs on 1 cpu, an instance of apache may be on the other cpu at the same time, answering some other query. So when you take in all the other stuff on a typical production mysql server, the multithreading in linux isn't going to make any difference for just 2 CPU's.

Ted

---

freebsd-questions@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-questions>

To unsubscribe, send any mail to "freebsd-questions-unsubscribe@freebsd.org"