

Re: Is the fsync() fake on FreeBSD6.1?

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-06/msg00052.html>

- *From:* Greg 'groggy' Lehey <grog@xxxxxxxxxxx>
 - *Date:* Tue, 27 Jun 2006 13:04:12 +0930
-

On Tuesday, 27 June 2006 at 10:18:47 +0800, leo huang wrote:

Hi,

I benchmarked MySQL 4.1.18 on FreeBSD 6.1 and Debian 3.1 using Super Smack 1.3 some days ago.

...

The result surprise me. The MySQL Performance on FreeBSD6.1 is about 10 times of on Debian3.1??and the output of iostat also shows it.

I know that MySQL uses fsync() to flush both the data and log files at default when using innodb engine(<http://dev.mysql.com/doc/refman/4.1/en/innodb-parameters.html>). Our evaluating computer only has a 10000RPM SCSI hard disk. I think it can do about 200 sequential fsync() calls per second if the fsync() is real.

Is the fsync() on FreeBSD6.1 fake?

My understanding from the last time I looked at the code was that fsync does the right thing:

The fsync() system call causes all modified data and attributes of fd to be moved to a permanent storage device. This normally results in all in-core modified copies of buffers for the associated file to be written to a disk.

This is not the case for Linux, where fsync syncs the entire file system. That could explain some of the performance difference, but not all of it. I suppose it's worth noting that, in general, people report much better performance with MySQL on Linux than on FreeBSD.

I mean than the data is only written to the drives memory and so can be lost if power goes down.

Re: Is the fsync() fake on FreeBSD6.1?

I don't believe that fsync is required to flush the drive buffers. It would be nice to have a function that did, though.

And how I can confirm this?

Trial and error?

Greg

—

See complete headers for address and phone numbers.

Attachment: pgpnNrK5hPCzf.pgp

Description: PGP signature