
Subject: Re: seems to be a flaw in cfq

Posted by [Jens Axboe](#) on Tue, 19 Dec 2006 18:47:48 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, Dec 19 2006, Jens Axboe wrote:

> On Tue, Dec 19 2006, Jens Axboe wrote:

> > On Tue, Dec 19 2006, Jens Axboe wrote:

> > > On Tue, Dec 19 2006, Vasily Tarasov wrote:

> > > > ello, Jens.

> > > >

> > > > Seems, that we've found some problem in CFQ.

> > > > I used your fio tool of version 1.9 to reproduce it.

> > > > 2.6.18 vanilla kernel.

> > >

> > > I'll look over this report - in the mean time, can you see if the same

> > > situation exists in 2.6.19 and 2.6.20-rc1? Would help a lot!

> >

> > I just tried to reproduce it in a recent kernel, and it does show

> > something very close to what you reported. Using 5 threads, 2 of them

> > sharing the same file, the 3 first threads get 6067 -> 6182KiB/sec each,

> > and the two threads sharing a file get 975KiB/sec each. That's really

> > bizarre, I'll take a good look at this! Thanks for reporting it.

>

> Oh, thinking about this - it could be an artifact of being too fair. The

> default is to use O_DIRECT, so the sharing threads probably end up being

> blocked waiting for each other to finish the same blocks of io. They

> will both be reading from the start of the file to the end, so if they

> run alongside each other inside the file, they'll be blocking each other

> waiting for io to finish. The same should happen for AS, though. Perhaps

> it's bad alias handling in CFQ.

>

> If I add offset=512m to the last thread so that the two sharing threads

> read different parts of the file, the result is completely fair

> (4.9MiB/sec -> 5.1MiB/sec for each thread).

>

> I'll keep looking.

Back after dinner, the fresh energy served it's purpose - I think I know what the issue is. We allow merging across process queues, which will effectively serialize some io if they are sync (like this case). I'll hack up a fix for current git and give it a test spin, to verify that this is the problem here.

--

Jens Axboe
