Subject: Re: [PATCH] cfq: get rid of cfqq hash
Posted by Jens Axboe on Wed, 25 Apr 2007 07:57:25 GMT
View Forum Message <> Reply to Message

On Wed, Apr 25 2007, Vasily Tarasov wrote:
> >> @@ -1806,7 +1765,11 @@ static int cfq_may_queue(request_queue_t
> >>   * so just lookup a possibly existing queue, or return 'may queue'
> >>   * if that fails
> >>   */
> >> - cfqq = cfq_find_cfq_hash(cfqd, key, tsk->ioprio);
> >> + cic = cfq_get_io_context_noalloc(cfqd, tsk);
> >> + if (!cic)
> >> + return ELV_MQUEUE_MAY;
> >> +
> >> + cfqq = cic->cfqq[rw & REQ_RW_SYNC];
> >>   if (cfqq) {
> >>   cfq_init_prio_data(cfqq);
> >>   cfq_prio_boost(cfqq);
> >
> > Ahem, how well did you test this patch?
>
> Ugh, again: bio_sync() returns not only 0/1
> Sorry for giving so much trouble...

Right, and REQ_RW_SYNC isn't 1 either, so it returns a large number if
set.

> BTW, what tests do you use to check patches?
> I'll run them on our nodes each time when sending it to you.
> At the moment I use some self made tests and a bit fio scripts.

I went to run a test testing many disks, with a fio file like so:

[root@AS4 ~]# cat many-rw-256
[global]
rw=write
bs=256k
direct=1
ioengine=libaio
iodepth=4096

[md0]
file_service_type=roundrobin:16
filename=/dev/sdix:/dev/sdiw:/dev/sdiv:...

filename is 256 scsi disks, using scsi_debug. I wanted to evaluate the
possible extra CPU usage from one process with a lot of io contexts
attached. And the benefits of such a patch as this one:

http://git.kernel.dk/?p=linux-2.6-block.git
;a=commitdiff;h=7e950c8181e63345743130d839680999c5de968a;hp=
551e9405cb9e1f900da456ba57ddcf35dea110b9

--
Jens Axboe
