Subject: Re: [PATCH] cfq: get rid of cfqq hash
Posted by Jens Axboe on Tue, 24 Apr 2007 17:02:28 GMT
View Forum Message <> Reply to Message

On Tue, Apr 24 2007, Vasily Tarasov wrote:
> From: Vasily Tarasov <vtaras@openvz.org>
>
> cfq hash is no more necessary. We always can get cfqq from io context.
> cfq_get_io_context_noalloc() function is introduced, because we don't want to
> allocate cic on merging and checking may_queue.
> In order to identify sync queue we've used hash key = CFQ_KEY_ASYNC. Since hash
> is eliminated we need to use other criterion: sync flag for queue is added.
> In all places where we dig in rb_tree we're in current context, so no
> additional locking is required.
>
> Advantages of this patch: no additional memory for hash, no seeking in hash,
> code is cleaner. But it is necessary now to seek cic in per-ioc rbtree, but
> it is faster:
> - most processes work only with few devices
> - most systems have only few block devices
> - it is a rb-tree

Vasily, this is still not against the CFQ branch, I get tons of rejects:

axboe@nelson:/src/linux-2.6-block $ patch -p1 --dry-run < ~/foo
[...]
10 out of 27 hunks FAILED -- saving rejects to file
block/cfq-iosched.c.rej

If you don't want to use the git tree, then just grab

http://brick.kernel.dk/snaps/cfq-update-20070424

and apply it to 2.6.21-rc7-gitX (latest) and provide a diff against
that. Thanks!

--
Jens Axboe