
Subject: Re: [PATCH] cfq: get rid of cfqq hash

Posted by [Vasily Tarasov](#) on Tue, 24 Apr 2007 12:51:06 GMT

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

>> 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, thanks for doing this, it's a good cleanup. One request, though

> - care to rebase this against the cfq devel branch? I'm afraid it's

> nowhere near applying right now.

> --

> Jens Axboe

Ouch, I see the sources are changed since the moment I was getting them...

I'll resend a fresh version!
