Subject: Re: [dm-devel] Re: dm: bounce_pfn limit added
Posted by Alasdair G Kergon on Thu, 01 Nov 2007 00:00:07 GMT
View Forum Message <> Reply to Message

On Wed, Oct 31, 2007 at 05:00:16PM -0500, Kiyoshi Ueda wrote:
> How about the case that other dm device is stacked on the dm device?
> (e.g. dm-linear over dm-multipath over i2o with bounce_pfn=64GB, and
> the multipath table is changed to i2o with bounce_pfn=1GB.)

Let's not broaden the problem out in that direction yet - that's a known flaw in the way all these device restrictions are handled. (Which would, it happens, also be resolved by the dm architectural changes I'm contemplating.)

Yes, we could certainly take this patch - it won't do much harm (just hit performance in some configurations). But I am not yet convinced that there isn't some further underlying problem with the way the responsibility for this bouncing is divided up between the various layers: I still don't feel I completely understand this problem yet.

- How does that bio_alloc() in blk_queue_bounce() guarantee never to lead a deadlock (in the device-mapper context)?
- Are some functions failing to take account of the hw_segments (and perhaps other) restrictions?
- Are things actually simpler if the bouncing is dealt with just once prior to entering the device stack (even though that may involve bouncing some data that does not need it) or is it better to endeavour to keep the bouncing as close to the final layer as possible?

Alasdair
--
agk@redhat.com