Subject: Re: [dm-devel] Re: dm: bounce_pfn limit added
Posted by Alasdair G Kergon on Wed, 31 Oct 2007 02:01:33 GMT
View Forum Message <> Reply to Message

On Tue, Oct 30, 2007 at 11:26:17PM +0000, Alasdair G Kergon wrote:
>) DM doesn't need to bounce bio's on its own, but the block layer defaults
>) to that in blk_queue_make_request(). The lower level drivers should
>) bounce ios themselves, that is what they need to do if not layered below
>) dm anyways.

So currently we treat bounce_pfn as a property that does not need to be
propagated through the stack.

But is that the right approach?
- Is there a blk_queue_bounce() missing either from dm or elsewhere?
  (And BTW can the bio_alloc() that lurks within lead to deadlock?)

Firstly, what's going wrong?
- What is the dm table you are using? (output of 'dmsetup table')
  - Which dm targets and with how many underlying devices?
- Which underlying driver?
- Is this direct I/O to the block device from userspace, or via some
filesystem or what?

Presumably you've shown the problem goes away when the dm device is
removed from the stack.
What if you swap in alternative dm targets, e.g. if it's linear,
try multipath (round-robin, one path)?

Alasdair
--
agk@redhat.com