
Subject: Re: [dm-devel] Re: dm: bounce_pfn limit added
Posted by [Alasdair G Kergon](#) on Tue, 30 Oct 2007 23:26:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Oct 30, 2007 at 01:11:38PM -0700, Andrew Morton wrote:

> On Mon, 29 Oct 2007 09:31:39 +0300

> Vasily Averin <vvs@sw.ru> wrote:

> > Device mapper uses its own bounce_pfn that may differ from one on underlying
> > device. In that way dm can build incorrect requests that contain sg elements
> > greater than underlying device is able to handle.

> > This is the cause of slab corruption in i2o layer, occurred on i386 arch when
> > very long direct IO requests are addressed to dm-over-i2o device.

Interesting. When this patch was originally proposed a couple of years ago, we did this instead:

<http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=commitdiff;h=daef265f1590cf3e6de989d074041a280c82d58b;hp=00d6da9b4d6707b808481372537adb0fb38f99b3>

) 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.

Alasdair

--

agk@redhat.com
