Subject: Re: [RFC][PATCH] EXT3: problem with page fault inside a transaction
Posted by Andrew Morton on Thu, 12 Oct 2006 06:43:30 GMT
View Forum Message <> Reply to Message

On Thu, 12 Oct 2006 09:57:26 +0400
Dmitriy Monakhov <dmonakhov@openvz.org> wrote:

> While reading Andrew's generic_file_buffered_write patches i've remembered
> one more EXT3 issue.journal_start() in prepare_write() causes different ranking
> violations if copy_from_user() triggers a page fault. It could cause
> GFP_FS allocation, re-entering into ext3 code possibly with a different
> superblock and journal, ranking violation of journalling serialization
> and mmap_sem and page lock and all other kinds of funny consequences.

With the stuff Nick and I are looking at, we won't take pagefaults inside
prepare_write()/commit_write() any more.

> Our customers complain about this issue.

Really? How often?

What on earth are they doing to trigger this? writev() without the 2.6.18
writev() bugfix?