

---

Subject: Re: [RFC][PATCH] EXT3: problem with page fault inside a transaction  
Posted by [Nick Piggin](#) on Thu, 12 Oct 2006 07:01:03 GMT

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

---

Andrew Morton wrote:

> 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.

Yep. Because the page is locked, it is too much to deal with even  
without a filesystem in the picture.

--

SUSE Labs, Novell Inc.

Send instant messages to your online friends <http://au.messenger.yahoo.com>

---