
Subject: Re: [RFC][PATCH] EXT3: problem with page fault inside a transaction
Posted by [Nick Piggin](#) on Thu, 12 Oct 2006 08:31:23 GMT

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

Dmitriy Monakhov wrote:

> Andrew Morton <akpm@osdl.org> writes:

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

>

> I'm sorry may be i've missed something, but how can't you prevent this?

>

> Let's look at generic_file_buffered_write:

> ##### force page fault

> fault_in_pages_readable();

>

> ### find and lock page

> __grab_cache_page()

>

> ##### allocate blocks. This may result in low memory condition

> ##### try_to_free_pages->shrink_caches() and etc.

> a_ops->prepare_write()

>

> ### can anyone guarantee that page fault hasn't happened by now ?

Yes. Do an atomic copy, which will early exit from the pagefault handler and return a short copy. Then close up the write, drop the page lock, and rerun the fault_in_pages_readable, which will do the full pagefaults for us, then try again.

Regardless of what you do to ext3, the VM just can't handle a fault here anyway.

--

SUSE Labs, Novell Inc.

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