

Re: help on lock around vm_page

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2007-06/msg00030.html>

- *From:* "Howard Su" <howard0su@xxxxxxxxxx>
 - *Date:* Thu, 7 Jun 2007 08:41:48 +0800
-

Some further reading give me a different message:

The pages returned by `vm_page_grab` should have `VPO_BUSY` flag which will prevent anyone else to move this page to cache. So current approach should be ok. Am I right? Maybe there is some edge cases since I only meet this bug in a low memory situation.

-Howard

On 6/7/07, Howard Su <howard0su@xxxxxxxxxx> wrote:

I want some helps from VM guru. I try to fix a panic in `tmpfs`. In order to push `tmpfs` into `-Current`, I really want some help to solve this.

1. we allocate an object from `vm_pager_alloc(OBJT_SWAP, ...)` when create a file.
2. the panic is during handling write op:
 - a) find the first page we want to write
 - b) call `vm_page_grab` to get the page from object.
 - c) call use `sf_buf_alloc` to map it into `kernel_map`
 - d) use `uiomove` to move the data
 - e) mark page as dirty
 - f) loop to a until all pages are handled.

there is a race condition. while doing b-c & e, we hold the `OBJ_LOCK/page_queue_lock`. when doing d, we have to drop the locks to call `uiomove`. When calling `uiomove`, the page may moved to cache queue since in that time it is not dirty.

There is a solution that we allocate a page buffer. Before a), we `uiomove` it to the buffer and replace `uiomove` with a `bcopy` in d). Then, we can hold lock in b - e. I feel this will cause performance problem.

For the detailed code, please check:

<http://perforce.freebsd.org/fileViewer.cgi?FSPC=//depot/user/howardsu/truss/sys/fs/tmpfs/tmpfs%5fvnops.c&>

Re: help on lock around vm_page

function: tmpfs_uio_xfer()

Any idea to close this race condition?

PS: If you can review my code about usage of vm, it will be appreciated.

--

-Howard

freebsd-arch@xxxxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-arch>

To unsubscribe, send any mail to "freebsd-arch-unsubscribe@xxxxxxxxxxxxx"