

Re: [PATCH] sendfile removal

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-05/msg13873.html>

- *From:* Jens Axboe <jens.axboe@xxxxxxxxxx>
 - *Date:* Thu, 31 May 2007 21:03:39 +0200
-

On Thu, May 31 2007, Hugh Dickins wrote:

On Thu, 31 May 2007, Jens Axboe wrote:

This patch removes the `->sendfile()` hook from the `file_operations` structure, and replaces the `sys_sendfile()` mechanism to be based on `->splice_read()` instead. There should be no functional changes.

Work to be done:

- The ext2 xip support needs a `splice_read()` implementation, currently I just if 0'ed out the `send_xip_file_sendfile()`. CC'ed Carsten, who seems to be the author of this code.
- `shmem` needs a `splice_read()` implementation. Optimistically CC'ed Hugh.

I'll take that "Optimistically" as a special compliment, rather than as a particular insult ;)

It's a compliment for sure, the "optimistically" was just mentioned because there are lots of people mentioned in that file, but I mostly remember you doing some heavy lifting there :-)

Yes, thanks, please leave `shmem_file_splice_read()` to me, I'll give it priority now. Not deep enough in yet, but I'll probably aim for something simple-minded (correct but slow once it hits swap).

Super, thanks!!

- `nfsd`: The `->rq_sendfile_ok` optimization is gone for now. I can't determine the value of it, but I'm assuming it's there for a reason. Any chance this can be converted to `splice`, or use something else than `->sendfile()`? CC'ed Neil.

Re: [PATCH] sendfile removal

– relay: needs a splice_read() implementation. I think Tom already has one, CC'ed him.

Apart from that, it was mostly straight forward. Almost everybody uses generic_file_sendfile(), which makes the conversion easy. I changed loop to use do_generic_file_read() instead of sendfile, it works for me...

Christoph already picked up on that, and it's of interest to shmem too: loop over tmpfs in 2.6 was relying on shmem_file_sendfile, for which the generic route is not good enough.

If we're giving a .splice_read to everything which used to have a .sendfile, then I think you just need to make do_lo_read() use ->splice_read now?

Yes, I will make that change myself. I can do that, as I seem to have sucessfully pushed out the remaining work to others :-)

--

Jens Axboe

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>