author	Jeff Layton <jlayton@kernel.org>	2023-03-17 13:13:08 -0400
committer	Greg Kroah-Hartman <gregkh@linuxfoundation.org>	2024-06-21 14:54:12 +0200
commit	8235cd619db6e67f1d7d26c55f1f3e4e575c947d (patch)	
tree	83cc835ed566b18ef834f7b9a45946e39fa833c2	
parent	37b34eb5677073ec6972f395dfe650cf44b421eb (diff)	
download	linux-8235cd619db6e67f1d7d26c55f1f3e4e575c947d.tar.gz	

nfsd: don't replace page in rq_pages if it's a continuation of last page

```
[ Upstream commit 27c934dd8832dd40fd34776f916dc201e18b319b ]

The splice read calls nfsd_splice_actor to put the pages containing file
data into the svc_rqst->rq_pages array. It's possible however to get a
splice result that only has a partial page at the end, if (e.g.) the
filesystem hands back a short read that doesn't cover the whole page.

nfsd_splice_actor will plop the partial page into its rq_pages array and
return. Then later, when nfsd_splice_actor is called again, the
remainder of the page may end up being filled out. At this point,
nfsd_splice_actor will put the page into the array _again_ corrupting
the reply. If this is done enough times, rq_next_page will overrun the
array and corrupt the trailing fields -- the rq_respages and
rq_next_page pointers themselves.

If we've already added the page to the array in the last pass, don't add
it to the array a second time when dealing with a splice continuation.
This was originally handled properly in nfsd_splice_actor, but commit
91e23b1c3982 ("NFSD: Clean up nfsd_splice_actor()") removed the check
for it.

Fixes: 91e23b1c3982 ("NFSD: Clean up nfsd_splice_actor()")
Cc: Al Viro <viro@zeniv.linux.org.uk>
Reported-by: Dario Lesca <d.lesca@solinos.it>
Tested-by: David Critch <dcritch@redhat.com>
Link: https://bugzilla.redhat.com/show_bug.cgi?id=2150630
Signed-off-by: Jeff Layton <jlayton@kernel.org>
Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
Signed-off-by: Sasha Levin <sashal@kernel.org>
```

Diffstat

-rw-r--r-- fs/nfsd/vfs.c 9

1 files changed, 8 insertions, 1 deletions

```
diff --git a/fs/nfsd/vfs.c b/fs/nfsd/vfs.c
index ddf424d76d4104..abc682854507ba 100644
--- a/fs/nfsd/vfs.c
+++ b/fs/nfsd/vfs.c
@@ -954,8 +954,15 @@ nfsd_splice_actor(struct pipe_inode_info *pipe, struct pipe_buffer *buf,
        struct page *last_page;
```

```
        last_page = page + (offset + sd->len - 1) / PAGE_SIZE;
-       for (page += offset / PAGE_SIZE; page <= last_page; page++)
+       for (page += offset / PAGE_SIZE; page <= last_page; page++) {
+               /*
+                * Skip page replacement when extending the contents
+                * of the current page.
+                */
+               if (page == *(rqstp->rq_next_page - 1))
+                       continue;
                svc_rqst_replace_page(rqstp, page);
+       }
        if (rqstp->rq_res.page_len == 0)        // first call
                rqstp->rq_res.page_base = offset % PAGE_SIZE;
        rqstp->rq_res.page_len += sd->len;
```