author	Jingbo Xu <jefflexu@linux.alibaba.com>	2022-11-14 20:19:43 +0800
committer	Gao Xiang <hsiangkao@linux.alibaba.com>	2022-11-14 23:48:38 +0800
commit	37020bbb71d911431e16c2c940b97cf86ae4f2f6 (patch)	
tree	4d62cb1eda9dd8b6ee2f1059a9746492412c11c4	
parent	39bfcb8138f6dc3375f23b1e62ccfc7c0d83295d (diff)	
download	linux-37020bbb71d911431e16c2c940b97cf86ae4f2f6.tar.gz	

erofs: fix missing xas_retry() in fscache mode

```
The xarray iteration only holds the RCU read lock and thus may encounter
XA_RETRY_ENTRY if there's process modifying the xarray concurrently.
This will cause oops when referring to the invalid entry.

Fix this by adding the missing xas_retry(), which will make the
iteration wind back to the root node if XA_RETRY_ENTRY is encountered.

Fixes: d435d53228dd ("erofs: change to use asynchronous io for fscache readpage/readahead")
Suggested-by: David Howells <dhowells@redhat.com>
Reviewed-by: Gao Xiang <hsiangkao@linux.alibaba.com>
Reviewed-by: Jia Zhu <zhujia.zj@bytedance.com>
Signed-off-by: Jingbo Xu <jefflexu@linux.alibaba.com>
Link: https://lore.kernel.org/r/20221114121943.29987-1-jefflexu@linux.alibaba.com
Signed-off-by: Gao Xiang <hsiangkao@linux.alibaba.com>
```

Diffstat

-rw-r--r-- fs/erofs/fscache.c 10

1 files changed, 7 insertions, 3 deletions

```
diff --git a/fs/erofs/fscache.c b/fs/erofs/fscache.c
index 6eaf4a4ab95ca1..af5ed6b9c54dd8 100644
--- a/fs/erofs/fscache.c
+++ b/fs/erofs/fscache.c
@@ -75,11 +75,15 @@ static void erofs_fscache_rreq_unlock_folios(struct netfs_io_request *rreq)

        rcu_read_lock();
        xas_for_each(&xas, folio, last_page) {
-               unsigned int pgpos =
-                       (folio_index(folio) - start_page) * PAGE_SIZE;
-               unsigned int pgend = pgpos + folio_size(folio);
+               unsigned int pgpos, pgend;
                bool pg_failed = false;

+               if (xas_retry(&xas, folio))
+                       continue;
+
+               pgpos = (folio_index(folio) - start_page) * PAGE_SIZE;
+               pgend = pgpos + folio_size(folio);
+
                for (;;) {
                        if (!subreq) {
```

```
pg_failed = true;
```