author	Baokun Li <libaokun1@huawei.com>	2023-01-07 11:21:26 +0800
committer	Greg Kroah-Hartman <gregkh@linuxfoundation.org>	2023-03-22 13:33:54 +0100
commit	70e66bdeae4d0f7c8e87762f425b68aedd5e8955 (patch)	
tree	8f5c2e5b654d1c7d36bfd4f018b0151ec6bead83	
parent	b4afe4629ec8f7d740617866e4286717a1c9c4c3 (diff)	
download	linux-70e66bdeae4d0f7c8e87762f425b68aedd5e8955.tar.gz	

ext4: update s_journal_inum if it changes after journal replay

```
[ Upstream commit 3039d8b8692408438a618fac2776b629852663c3 ]

When mounting a crafted ext4 image, s_journal_inum may change after journal
replay, which is obviously unreasonable because we have successfully loaded
and replayed the journal through the old s_journal_inum. And the new
s_journal_inum bypasses some of the checks in ext4_get_journal(), which
may trigger a null pointer dereference problem. So if s_journal_inum
changes after the journal replay, we ignore the change, and rewrite the
current journal_inum to the superblock.

Link: https://bugzilla.kernel.org/show_bug.cgi?id=216541
Reported-by: Luís Henriques <lhenriques@suse.de>
Signed-off-by: Baokun Li <libaokun1@huawei.com>
Reviewed-by: Jan Kara <jack@suse.cz>
Link: https://lore.kernel.org/r/20230107032126.4165860-3-libaokun1@huawei.com
Signed-off-by: Theodore Ts'o <tytso@mit.edu>
Signed-off-by: Sasha Levin <sashal@kernel.org>
```

Diffstat

-rw-r--r-- fs/ext4/super.c 7

1 files changed, 5 insertions, 2 deletions

```
diff --git a/fs/ext4/super.c b/fs/ext4/super.c
index 80116009995864..2528e8216c3342 100644
--- a/fs/ext4/super.c
+++ b/fs/ext4/super.c
@@ -5967,8 +5967,11 @@ static int ext4_load_journal(struct super_block *sb,
        if (!really_read_only && journal_devnum &&
            journal_devnum != le32_to_cpu(es->s_journal_dev)) {
                es->s_journal_dev = cpu_to_le32(journal_devnum);
-
-               /* Make sure we flush the recovery flag to disk. */
+               ext4_commit_super(sb);
+       }
+       if (!really_read_only && journal_inum &&
+           journal_inum != le32_to_cpu(es->s_journal_inum)) {
+               es->s_journal_inum = cpu_to_le32(journal_inum);
                ext4_commit_super(sb);
        }
```