author	Wang Yufen <wangyufen@huawei.com>	2022-11-02 17:41:19 +0800
committer	Greg Kroah-Hartman <gregkh@linuxfoundation.org>	2022-11-16 09:57:09 +0100
commit	3401f964028ac941425b9b2c8ff8a022539ef44a (patch)	
tree	9b96e14229b07a0d7f7b470b828a46b7c51a9c11	
parent	adaa0f180de5236e086ddab6476c4364d922f1fd (diff)	
download	linux-3401f964028ac941425b9b2c8ff8a022539ef44a.tar.gz	

net: tun: Fix memory leaks of napi_get_frags

```
[ Upstream commit 1118b2049d77ca0b505775fc1a8d1909cf19a7ec ]

kmemleak reports after running test_progs:

unreferenced object 0xffff8881b1672dc0 (size 232):
  comm "test_progs", pid 394388, jiffies 4354712116 (age 841.975s)
  hex dump (first 32 bytes):
    e0 84 d7 a8 81 88 ff ff 80 2c 67 b1 81 88 ff ff  .........,g.....
    00 40 c5 9b 81 88 ff ff 00 00 00 00 00 00 00 00  .@..............
  backtrace:
    [<00000000c8f01748>] napi_skb_cache_get+0xd4/0x150
    [<0000000041c7fc09>] __napi_build_skb+0x15/0x50
    [<00000000431c7079>] __napi_alloc_skb+0x26e/0x540
    [<000000003ecfa30e>] napi_get_frags+0x59/0x140
    [<0000000099b2199e>] tun_get_user+0x183d/0x3bb0 [tun]
    [<000000008a5adef0>] tun_chr_write_iter+0xc0/0x1b1 [tun]
    [<0000000049993ff4>] do_iter_readv_writev+0x19f/0x320
    [<000000008f338ea2>] do_iter_write+0x135/0x630
    [<000000008a3377a4>] vfs_writev+0x12e/0x440
    [<00000000a6b5639a>] do_writev+0x104/0x280
    [<00000000ccf065d8>] do_syscall_64+0x3b/0x90
    [<00000000d776e329>] entry_SYSCALL_64_after_hwframe+0x63/0xcd

The issue occurs in the following scenarios:
tun_get_user()
  napi_gro_frags()
    napi_frags_finish()
      case GRO_NORMAL:
        gro_normal_one()
          list_add_tail(&skb->list, &napi->rx_list);
          <-- While napi->rx_count < READ_ONCE(gro_normal_batch),
          <-- gro_normal_list() is not called, napi->rx_list is not empty
  <-- not ask to complete the gro work, will cause memory leaks in
  <-- following tun_napi_del()
...
tun_napi_del()
  netif_napi_del()
    __netif_napi_del()
    <-- &napi->rx_list is not empty, which caused memory leaks

To fix, add napi_complete() after napi_gro_frags().
```

```
Fixes: 90e33d459407 ("tun: enable napi_gro_frags() for TUN/TAP driver")
Signed-off-by: Wang Yufen <wangyufen@huawei.com>
Reviewed-by: Eric Dumazet <edumazet@google.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
Signed-off-by: Sasha Levin <sashal@kernel.org>
```

Diffstat

-rw-r--r-- drivers/net/tun.c 1

1 files changed, 1 insertions, 0 deletions

```
diff --git a/drivers/net/tun.c b/drivers/net/tun.c
index 0c09f8e9d38365..83662f616b679b 100644
--- a/drivers/net/tun.c
+++ b/drivers/net/tun.c
@@ -1996,6 +1996,7 @@ drop:

                local_bh_disable();
                napi_gro_frags(&tfile->napi);
+               napi_complete(&tfile->napi);
                local_bh_enable();
                mutex_unlock(&tfile->napi_mutex);
        } else if (tfile->napi_enabled) {
```