

---

Subject: [PATCH] Don't attach callback to a going-away netlink socket

Posted by [xemul](#) on Mon, 16 Apr 2007 11:34:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

From: Denis Lunev <den@openvz.org>

There is a race between netlink\_dump\_start() and netlink\_release() that can lead to the situation when a netlink socket with non-zero callback is freed.

```
--- a/net/netlink/af_netlink.c 2004-10-25 12:12:23.000000000 +0400
+++ b/net/netlink/af_netlink.c 2004-10-28 16:26:12.000000000 +0400
@@ -255,6 +255,7 @@ static int netlink_release(struct socket
     return 0;

     netlink_remove(sk);
+ sock_orphan(sk);
     nlk = nlk_sk(sk);

     spin_lock(&nlk->cb_lock);
@@ -269,7 +270,6 @@ static int netlink_release(struct socket
     /* OK. Socket is unlinked, and, therefore,
        no new packets will arrive */

- sock_orphan(sk);
     sock->sk = NULL;
     wake_up_interruptible_all(&nlk->wait);

@@ -942,9 +942,9 @@ int netlink_dump_start(struct sock *ssk,
     return -ECONNREFUSED;
 }
     nlk = nlk_sk(sk);
- /* A dump is in progress... */
+ /* A dump or destruction is in progress... */
     spin_lock(&nlk->cb_lock);
- if (nlk->cb) {
+ if (nlk->cb || sock_flag(sk, SOCK_DEAD)) {
     spin_unlock(&nlk->cb_lock);
     netlink_destroy_callback(cb);
     sock_put(sk);
```

---