Subject: [PATCH] [NETNS 1/4 net-2.6.25] Double free in netlink_release.
Posted by den on Fri, 18 Jan 2008 12:53:13 GMT
View Forum Message <> Reply to Message

Netlink protocol table is global for all namespaces. Some netlink protocols
have been virtualized, i.e. they have per/namespace netlink socket. This
difference can easily lead to double free if more than 1 namespace is
started. Count the number of kernel netlink sockets to track that this
table is not used any more.

Signed-off-by: Denis V. Lunev <den@openvz.org>
Tested-by: Alexey Dobriyan <adobriyan@openvz.org>
---
 net/netlink/af_netlink.c | 10 +++++++---
 1 files changed, 7 insertions(+), 3 deletions(-)

diff --git a/net/netlink/af_netlink.c b/net/netlink/af_netlink.c
index 21f9e30..29fef55 100644
--- a/net/netlink/af_netlink.c
+++ b/net/netlink/af_netlink.c
@@ -498,9 +498,12 @@ static int netlink_release(struct socket *sock)

  netlink_table_grab();
  if (netlink_is_kernel(sk)) {
- kfree(nl_table[sk->sk_protocol].listeners);
- nl_table[sk->sk_protocol].module = NULL;
- nl_table[sk->sk_protocol].registered = 0;
+ BUG_ON(nl_table[sk->sk_protocol].registered == 0);
+ if (--nl_table[sk->sk_protocol].registered == 0) {
+   kfree(nl_table[sk->sk_protocol].listeners);
+   nl_table[sk->sk_protocol].module = NULL;
+   nl_table[sk->sk_protocol].registered = 0;
+ }
  } else if (nlk->subscriptions)
   netlink_update_listeners(sk);
  netlink_table_ungrab();
@@ -1389,6 +1392,7 @@ netlink_kernel_create(struct net *net, int unit, unsigned int groups,
   nl_table[unit].registered = 1;
  } else {
   kfree(listeners);
+ nl_table[unit].registered++;
  }
  netlink_table_ungrab();

--
1.5.3.rc5