
Subject: [patch 1/3][netns] add a reference to the netns for timewait

Posted by [Daniel Lezcano](#) on Mon, 24 Sep 2007 13:29:36 GMT

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

From: Daniel Lezcano <dlezcano@fr.ibm.com>

When a socket changes to a timewait socket, the network namespace is not copied from the original socket.

Here we hold a usage reference, not the ref count on the network namespace, so the network namespace will be freed either the usage reference is not 0. The network namespace cleanup function will fail if there is any usage of it. In this case, we should ensure there is no usage of the network namespace.

Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>

include/net/inet_timewait_sock.h | 2 ++
net/ipv4/inet_timewait_sock.c | 1 +
2 files changed, 3 insertions(+)

Index: linux-2.6-netns/include/net/inet_timewait_sock.h

```
=====
--- linux-2.6-netns.orig/include/net/inet_timewait_sock.h
+++ linux-2.6-netns/include/net/inet_timewait_sock.h
@@ -197,12 +197,14 @@ static inline void inet_tws_k_put(struct
 {
     if (atomic_dec_and_test(&tw->tw_refcnt)) {
         struct module *owner = tw->tw_prot->owner;
+ struct net *net = tw->tw_net;
         twsk_destructor((struct sock *)tw);
#ifdef SOCK_REFCNT_DEBUG
         printk(KERN_DEBUG "%s timewait_sock %p released\n",
             tw->tw_prot->name, tw);
#endif
         kmem_cache_free(tw->tw_prot->tws_k_prot->tws_k_slab, tw);
+ release_net(net);
         module_put(owner);
     }
 }
```

Index: linux-2.6-netns/net/ipv4/inet_timewait_sock.c

```
=====
--- linux-2.6-netns.orig/net/ipv4/inet_timewait_sock.c
+++ linux-2.6-netns/net/ipv4/inet_timewait_sock.c
@@ -108,6 +108,7 @@ struct inet_timewait_sock *inet_tws_k_all
     tw->tw_hash = sk->sk_hash;
     tw->tw_ipv6only = 0;
     tw->tw_prot = sk->sk_prot_creator;
```

```
+ tw->tw_net      = hold_net(sk->sk_net);  
  atomic_set(&tw->tw_refcnt, 1);  
  inet_twsk_dead_node_init(tw);  
  __module_get(tw->tw_prot->owner);
```

--

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
