
Subject: [PATCH 2/8 net-2.6.26] [NETNS]: Add netns refcnt debug for kernel sockets.

Posted by [den](#) on Tue, 15 Apr 2008 12:37:49 GMT

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

Protocol control sockets and netlink kernel sockets should not prevent the namespace stop request. They are initialized and disposed in a special way by `sk_change_net/sk_release_kernel`.

Signed-off-by: Denis V. Lunev <den@openvz.org>

```
include/net/sock.h | 2 +-
net/core/sock.c   | 1 +
2 files changed, 2 insertions(+), 1 deletions(-)
```

```
diff --git a/include/net/sock.h b/include/net/sock.h
```

```
index 09255ea..dc42b44 100644
```

```
--- a/include/net/sock.h
```

```
+++ b/include/net/sock.h
```

```
@@ -1314,7 +1314,7 @@ void sock_net_set(struct sock *sk, struct net *net)
```

```
static inline void sk_change_net(struct sock *sk, struct net *net)
```

```
{
    put_net(sock_net(sk));
- sock_net_set(sk, net);
+ sock_net_set(sk, hold_net(net));
}
```

```
extern void sock_enable_timestamp(struct sock *sk);
```

```
diff --git a/net/core/sock.c b/net/core/sock.c
```

```
index f2ccb16..015ec69 100644
```

```
--- a/net/core/sock.c
```

```
+++ b/net/core/sock.c
```

```
@@ -1001,6 +1001,7 @@ void sk_release_kernel(struct sock *sk)
```

```
    sock_hold(sk);
    sock_release(sk->sk_socket);
+ release_net(sock_net(sk));
    sock_net_set(sk, get_net(&init_net));
    sock_put(sk);
}
```

```
--
```

```
1.5.3.rc5
```

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
