
Subject: Re: [PATCH 06/16] net: Add a network namespace parameter to struct sock

Posted by [Daniel Lezcano](#) on Thu, 20 Sep 2007 13:20:01 GMT

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

Denis V. Lunev wrote:

> Eric W. Biederman wrote:

>> Sockets need to get a reference to their network namespace,
>> or possibly a simple hold if someone registers on the network
>> namespace notifier and will free the sockets when the namespace
>> is going to be destroyed.

>>

>> Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

>> ---

>> include/net/inet_timewait_sock.h | 1 +

>> include/net/sock.h | 3 +++

>> 2 files changed, 4 insertions(+), 0 deletions(-)

>>

>> diff --git a/include/net/inet_timewait_sock.h b/include/net/inet_timewait_sock.h

>> index 47d52b2..abaff05 100644

>> --- a/include/net/inet_timewait_sock.h

>> +++ b/include/net/inet_timewait_sock.h

>> @@ -115,6 +115,7 @@ struct inet_timewait_sock {

>> #define tw_refcnt __tw_common.skc_refcnt

>> #define tw_hash __tw_common.skc_hash

>> #define tw_prot __tw_common.skc_prot

>> +#define tw_net __tw_common.skc_net

>

>

> This place is a very tricky, indeed. If we keep the namespace until

> timewait bucket death - we'll keep the namespace alive at least 5

> _minutes_ after all process death.

Yes, that's right. And for me that makes totally sense. The namespace should not be destroyed until it is referenced somewhere.

> If we stop a VE (in terms of OpenVz) and restart it, we'll 100% have an

> _OLD_ namespace with all buckets shown :(

> So, in OpenVz we use a number

> of VE instead of pointer to a VE. Additionally, on VE death we can wipe

> all TW buckets. VE start stop from outside world looks very much like a

> computer power on/off.

That makes sense too. But if you wipe out the sockets when stopping the VE where is the problem with the restart ?
