
Subject: Re: [PATCH][NETNS] Use list_for_each_entry_continue_reverse in setup_net

Posted by [Stephen Hemminger](#) on Fri, 14 Sep 2007 20:06:56 GMT

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

On Fri, 14 Sep 2007 08:41:07 -0600

ebiederm@xmission.com (Eric W. Biederman) wrote:

> Stephen Hemminger <shemminger@linux-foundation.org> writes:

>

>> On Fri, 14 Sep 2007 11:39:32 +0400

>> Pavel Emelyanov <xemul@openvz.org> wrote:

>>

>>> I proposed introducing a list_for_each_entry_continue_reverse

>>> macro to be used in setup_net() when unrolling the failed

>>> ->init callback.

>>>

>>> Here is the macro and some more cleanup in the setup_net() itself

>>> to remove one variable from the stack :) Minor, but the code

>>> looks nicer.

>>>

>>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

>>>

>>> Maybe it is time to just eliminate the init hook from the API.

>>> It has very few users, and there is no reason the setup needed

>>> could be done before or after registering in most cases.

>>>

>>> I guess only have 5 out of the 29 users I have in my full patchset

>>> is few. But that is to be expected because so far only the core

>>> has been converted.

>>>

>>> I looked again at the initialization to see if you had a point about

>>> the initialization but in every instance I looked at the function

>>> was performing work that needed to happen during the creation of

>>> each network namespace. So the work very much needs to be done there.

>>>

>>> Ok looking some more I can see why this isn't obvious yet. copy_net_ns

>>> hasn't been merged yet, and that is where we create new network namespaces.

>>> And call setup_net on each new network namespace.

>>>

>>> I will take a look at that patch and see if I can come up with a

>>> safe version of it to merge to allow for a little more transparency.

Could we just make it so dev->init is not allowed to fail? Then it can be a void function and the nasty unwind code can go?
