Subject: Re: [patch 1/2][NETNS][RFD] store the network namespace pointer in the dst_entry structure
Posted by ebiederm on Tue, 11 Dec 2007 15:52:24 GMT
View Forum Message <> Reply to Message

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

> Store the network namespace pointer in the dst_entry structure when it is
> allocated.
> The different protocols redefine the route object as a derivate object from
> dst_entry. So using the dst_entry to store the network namespace pointer will
> allow to take into account the ipv4, ipv6, dccp protocols in one shot through
> the different route objects, rtable, rt6_info, ...
>
> ---
> include/net/dst.h | 3 ++-
> net/core/dst.c | 3 ++-
> net/decnet/dn_route.c | 4 ++--
> net/ipv4/route.c | 14 +++++++-------
> net/ipv6/route.c | 18 ++++++++++--------
> net/xfrm/xfrm_policy.c | 2 +-
> 6 files changed, 24 insertions(+), 20 deletions(-)
>
> Index: linux-2.6-netns/include/net/dst.h
> ===
> --- linux-2.6-netns.orig/include/net/dst.h
> +++ linux-2.6-netns/include/net/dst.h
> @@ -81,6 +81,7 @@ struct dst_entry
>   struct dn_route *dn_next;
>   };
>   char   info[0];
> + struct net *net;

Unless I'm missing something you just place that net pointer in
the middle of a variable length array. Weird I don't see us
using that array.

Could you please place the struct net *net pointer up by the
network device pointer.
> };

I know we need a net pointer in struct rt_table, because it
is a hash table that we can't dynamically allocate so we need
to place a network namespace pointer as part of the hash key.

For the ipv6 fib tables I don't recall needing a net pointer
as we didn't have a hash table and could instead have separate
roots for different namespaces.

I find this slightly odd as I didn't wind up needing to add
a struct net pointer in struct dst in my proof of concept tree
and struct dst doesn't have a struct flowi so that would not
have prevented it.

Eric

_____

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