
Subject: Re: [PATCH 6/17 net-2.6.26] [NETNS]: Default arp parameters lookup.
Posted by [Daniel Lezcano](#) on Tue, 19 Feb 2008 15:16:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daniel Lezcano wrote:

> Denis V. Lunev wrote:

>> On Tue, 2008-02-19 at 10:51 +0100, Daniel Lezcano wrote:

>>> Denis V. Lunev wrote:

>>>> On Tue, 2008-02-19 at 10:14 +0100, Daniel Lezcano wrote:

>>>>> Denis V. Lunev wrote:

>>>>>> Default ARP parameters should be findable regardless of the context.

>>>>>> Required to make inetdev_event working.

>>>>>>

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

>>>>>> ---

>>>>>> net/core/neighbour.c | 4 +---

>>>>>> 1 files changed, 1 insertions(+), 3 deletions(-)

>>>>>>

>>>>>> diff --git a/net/core/neighbour.c b/net/core/neighbour.c

>>>>>> index c895ad4..45ed620 100644

>>>>>> --- a/net/core/neighbour.c

>>>>>> +++ b/net/core/neighbour.c

>>>>>> @@ -1275,9 +1275,7 @@ static inline struct neigh_parms

>>>>>> *lookup_neigh_params(struct neigh_table *tbl,

>>>>>> struct neigh_parms *p;

>>>>>>

>>>>>> for (p = &tbl->parms; p; p = p->next) {

>>>>>> - if (p->net != net)

>>>>>> - continue;

>>>>>> - if ((p->dev && p->dev->ifindex == ifindex) ||

>>>>>> + if ((p->dev && p->dev->ifindex == ifindex && p->net ==

>>>>>> net) ||

>>>>>> (!p->dev && !ifindex))

>>>>>> return p;

>>>>>> }

>>>>> If the values are:

>>>>> p->dev == NULL

>>>>> ifindex == 0

>>>>> p->net != net

>>>>>

>>>>> The parms should not be taken into account and the looping must

>>>>> continue. But with this modification it is not the case, if we

>>>>> specify parms ifindex == 0, the first parms with the dev field set

>>>>> to NULL will be taken belonging or not to the right net.

>>>>> They should be taken. In the other case inetdev_event will fail for

>>>>> sure

>>>>> in the middle. You could check.

>>>>>

>>>> These are ARP defaults and I do not see a problem for now to get them.
>>> Because there is a parms default per namespace. So several instances
>>> of them per nd table. That was the initial approach with Eric's
>>> patchset.
>>>
>>
>> These changes are not in mainstream and I do not want to touch ARP as
>> this is not a simple thing. In reality ARP will be needed only when
>> we'll have a real device inside a namespace.
>>
>> Right now I prefer to have minimal set of working changes to finish IP
>> and upper layers.
>
> core/neighbour.c is a common part between several protocols, especially
> ipv4 and ipv6. If you modify this function just to fit your need in the
> arp that will block me for ipv6 until you make parms default per
> namespace. So please, find another way to do that, perhaps just add a
> helper function.
>
> I suggest you do parms default per namespace first, it is quite small
> and easy :)
>
> Just let me the time to send the copy-parms-default function.

Ok, so after a long discussion with Denis about this patch, I will
change the ipv6 code to share the neigh->parms. It is not a problem.
Having the behavior of the neighbour subsystem per namespace is not a
must-have.

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