
Subject: [PATCH 5/6 net-2.6.25] [NETNS] Pass correct namespace in ip_route_input_slow.

Posted by [den](#) on Mon, 21 Jan 2008 14:50:45 GMT

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

The packet on the input path always has a reference to an input network device it is passed from. Extract network namespace from it.

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

net/ipv4/route.c | 7 +++++-

1 files changed, 4 insertions(+), 3 deletions(-)

diff --git a/net/ipv4/route.c b/net/ipv4/route.c

index c107bc3..b3c6122 100644

--- a/net/ipv4/route.c

+++ b/net/ipv4/route.c

@@ -1881,6 +1881,7 @@ static int ip_route_input_slow(struct sk_buff *skb, __be32 daddr,
__be32 saddr,
__be32 spec_dst;
int err = -EINVAL;
int free_res = 0;
+ struct net * net = dev->nd_net;

/* IP on this device is disabled. */

@@ -1911,7 +1912,7 @@ static int ip_route_input_slow(struct sk_buff *skb, __be32 daddr,
__be32 saddr,

/*

* Now we are ready to route packet.

*/

- if ((err = fib_lookup(&init_net, &fl, &res)) != 0) {

+ if ((err = fib_lookup(net, &fl, &res)) != 0) {

if (!IN_DEV_FORWARD(in_dev))

goto e_hostunreach;

goto no_route;

@@ -1926,7 +1927,7 @@ static int ip_route_input_slow(struct sk_buff *skb, __be32 daddr,
__be32 saddr,

if (res.type == RTN_LOCAL) {

int result;

result = fib_validate_source(saddr, daddr, tos,

- init_net.loopback_dev->ifindex,

+ net->loopback_dev->ifindex,

dev, &spec_dst, &itag);

if (result < 0)

goto martian_source;

@@ -1988,7 +1989,7 @@ local_input:

#endif

```
rth->rt_iif =  
rth->fl.iif = dev->ifindex;  
- rth->u.dst.dev = init_net.loopback_dev;  
+ rth->u.dst.dev = net->loopback_dev;  
dev_hold(rth->u.dst.dev);  
rth->idev = in_dev_get(rth->u.dst.dev);  
rth->rt_gateway = daddr;  
--
```

1.5.3.rc5
