
Subject: [PATCH 9/11 net-2.6.26] [NETNS]: Drop packets in the non-initial namespace on the per/protocol basis

Posted by [den](#) on Mon, 24 Mar 2008 14:36:10 GMT

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

IP layer now can handle multiple namespaces normally. So, process such packets normally and drop them only if the transport layer is not aware about namespaces.

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

```
include/net/protocol.h | 3 +-
net/ipv4/ip_input.c    | 8 ++++----
2 files changed, 6 insertions(+), 5 deletions(-)
```

```
diff --git a/include/net/protocol.h b/include/net/protocol.h
```

```
index ad8c584..8d024d7 100644
```

```
--- a/include/net/protocol.h
```

```
+++ b/include/net/protocol.h
```

```
@@ -39,7 +39,8 @@ struct net_protocol {
    int (*gso_send_check)(struct sk_buff *skb);
    struct sk_buff      (*gso_segment)(struct sk_buff *skb,
        int features);
- int no_policy;
+ unsigned int no_policy:1,
+ netns_ok:1;
};
```

```
#if defined(CONFIG_IPV6) || defined (CONFIG_IPV6_MODULE)
```

```
diff --git a/net/ipv4/ip_input.c b/net/ipv4/ip_input.c
```

```
index eb1fa27..2aeea5d 100644
```

```
--- a/net/ipv4/ip_input.c
```

```
+++ b/net/ipv4/ip_input.c
```

```
@@ -199,6 +199,8 @@ int ip_call_ra_chain(struct sk_buff *skb)
```

```
static int ip_local_deliver_finish(struct sk_buff *skb)
```

```
{
+ struct net *net = skb->dev->nd_net;
+
    __skb_pull(skb, ip_hdrlen(skb));
```

```
/* Point into the IP datagram, just past the header. */
```

```
@@ -214,7 +216,8 @@ static int ip_local_deliver_finish(struct sk_buff *skb)
    raw = raw_local_deliver(skb, protocol);
```

```
    hash = protocol & (MAX_INET_PROTOS - 1);
- if ((ipprot = rcu_dereference(inet_protos[hash])) != NULL) {
+ ipprot = rcu_dereference(inet_protos[hash]);
```

```
+ if (ipprot != NULL && (net == &init_net || ipprot->netns_ok)) {
    int ret;

    if (!ipprot->no_policy) {
@@ -375,9 +378,6 @@ int ip_rcv(struct sk_buff *skb, struct net_device *dev, struct packet_type
*pt,
    struct iphdr *iph;
    u32 len;

- if (dev->nd_net != &init_net)
- goto drop;
-
/* When the interface is in promisc. mode, drop all the crap
* that it receives, do not try to analyse it.
*/
--
1.5.3.rc5
```
