Subject: unlock iptables in netns
Posted by Alexey Dobriyan on Tue, 10 Jun 2008 17:27:20 GMT
View Forum Message <> Reply to Message

Hi,

Den basically banned iptables in netns via this patch

```
--- a/net/netfilter/core.c
+++ b/net/netfilter/core.c
@@ -165,14 +165,6 @@ int nf_hook_slow(int pf, unsigned int hook, struct sk_buff *skb,
  unsigned int verdict;
  int ret = 0;

-#ifdef CONFIG_NET_NS
- struct net *net;
-
- net = indev == NULL ? dev_net(outdev) : dev_net(indev);
- if (net != &init_net)
-  return 1;
-#endif
-
  /* We may already have this, but read-locks nest anyway */
  rcu_read_lock();

--- a/net/netfilter/nf_sockopt.c
+++ b/net/netfilter/nf_sockopt.c
@@ -65,9 +65,6 @@ static struct nf_sockopt_ops *nf_sockopt_find(struct sock *sk, int pf,
 {
  struct nf_sockopt_ops *ops;

- if (sock_net(sk) != &init_net)
-  return ERR_PTR(-ENOPROTOOPT);
-
  if (mutex_lock_interruptible(&nf_sockopt_mutex) != 0)
   return ERR_PTR(-EINTR);
```

, however, at least some of netfilter pieces are ready for usage in netns
and it would be nice to unlock them before release.

If I'm deciphering chengelog correctly it's all about code which does
nf_register_hook{,s} but not netns-ready itself:

 br_netfilter.c
 iptable_mangle (via ip_route_me_harder)
 conntracking (both IPv4 and IPv6)
 NAT
 arptable_filter

selinux
decnet
ebtable_filter
ebtable_nat
ipt_CLUSTERIP

Patch above can be applied and we can mark above list as "depends !NET_NS"
and move on.

Comments? Den, was there something else you're afraid of?
