
Subject: [PATCH 4/12] L2 network namespace: devinet sysctl's checks

Posted by [Mishin Dmitry](#) on Wed, 06 Dec 2006 22:26:18 GMT

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

This patch prohibits sysctl's registrations for devices, initialized in non-init network namespace. This is required due to absence of the sysctl virtualization.

Signed-off-by: Dmitry Mishin <dim@openvz.org>

```
net/ipv4/devinet.c | 26 ++++++-----
1 file changed, 16 insertions(+), 10 deletions(-)
```

--- linux-2.6.19-rc6-mm2.orig/net/ipv4/devinet.c

+++ linux-2.6.19-rc6-mm2/net/ipv4/devinet.c

```
@@ -161,8 +161,9 @@ struct in_device *inetdev_init(struct ne
```

```
/* Reference in_dev->dev */
```

```
dev_hold(dev);
```

```
#ifdef CONFIG_SYSCTL
```

```
- neigh_sysctl_register(dev, in_dev->arp_parms, NET_IPV4,
```

```
- NET_IPV4_NEIGH, "ipv4", NULL, NULL);
```

```
+ if (current_net_ns == &init_net_ns)
```

```
+ neigh_sysctl_register(dev, in_dev->arp_parms, NET_IPV4,
```

```
+ NET_IPV4_NEIGH, "ipv4", NULL, NULL);
```

```
#endif
```

```
/* Account for reference dev->ip_ptr */
```

```
@@ -170,7 +171,8 @@ struct in_device *inetdev_init(struct ne
```

```
rcu_assign_pointer(dev->ip_ptr, in_dev);
```

```
#ifdef CONFIG_SYSCTL
```

```
- devinet_sysctl_register(in_dev, &in_dev->cnf);
```

```
+ if (current_net_ns == &init_net_ns)
```

```
+ devinet_sysctl_register(in_dev, &in_dev->cnf);
```

```
#endif
```

```
ip_mc_init_dev(in_dev);
```

```
if (dev->flags & IFF_UP)
```

```
@@ -210,13 +212,15 @@ static void inetdev_destroy(struct in_de
```

```
}
```

```
#ifdef CONFIG_SYSCTL
```

```
- devinet_sysctl_unregister(&in_dev->cnf);
```

```
+ if (current_net_ns == &init_net_ns)
```

```
+ devinet_sysctl_unregister(&in_dev->cnf);
```

```
#endif
```

```
dev->ip_ptr = NULL;
```

```

#ifdef CONFIG_SYSCTL
- neigh_sysctl_unregister(in_dev->arp_parms);
+ if (current_net_ns == &init_net_ns)
+ neigh_sysctl_unregister(in_dev->arp_parms);
#endif
    neigh_parms_release(&arp_tbl, in_dev->arp_parms);
    arp_ifdown(dev);
@@ -1104,11 +1108,13 @@ static int inetdev_event(struct notifier
    inetdev_changename(dev, in_dev);

#ifdef CONFIG_SYSCTL
- devinet_sysctl_unregister(&in_dev->cnf);
- neigh_sysctl_unregister(in_dev->arp_parms);
- neigh_sysctl_register(dev, in_dev->arp_parms, NET_IPV4,
-     NET_IPV4_NEIGH, "ipv4", NULL, NULL);
- devinet_sysctl_register(in_dev, &in_dev->cnf);
+ if (current_net_ns == &init_net_ns) {
+ devinet_sysctl_unregister(&in_dev->cnf);
+ neigh_sysctl_unregister(in_dev->arp_parms);
+ neigh_sysctl_register(dev, in_dev->arp_parms, NET_IPV4,
+     NET_IPV4_NEIGH, "ipv4", NULL, NULL);
+ devinet_sysctl_register(in_dev, &in_dev->cnf);
+ }
#endif
    break;
}

```

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
