
Subject: [PATCH 14/17 net-2.6.26] [NETNS]: Process devinet ioctl in the correct namespace.

Posted by [den](#) on Tue, 19 Feb 2008 07:31:34 GMT

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

Add namespace parameter to devinet_ioctl and locate device inside it for state changes.

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

```
include/linux/inetdevice.h | 2 +-
net/ipv4/af_inet.c         | 7 ++++---
net/ipv4/devinet.c         | 6 ++++---
net/ipv4/ipconfig.c        | 2 +-
4 files changed, 9 insertions(+), 8 deletions(-)
```

```
diff --git a/include/linux/inetdevice.h b/include/linux/inetdevice.h
```

```
index fc4e3db..da05ab4 100644
```

```
--- a/include/linux/inetdevice.h
```

```
+++ b/include/linux/inetdevice.h
```

```
@@ -129,7 +129,7 @@ extern int unregister_inetaddr_notifier(struct notifier_block *nb);
```

```
extern struct net_device *ip_dev_find(struct net *net, __be32 addr);
extern int inet_addr_onlink(struct in_device *in_dev, __be32 a, __be32 b);
-extern int devinet_ioctl(unsigned int cmd, void __user *);
+extern int devinet_ioctl(struct net *net, unsigned int cmd, void __user *);
extern void devinet_init(void);
extern struct in_device *inetdev_by_index(struct net *, int);
extern __be32 inet_select_addr(const struct net_device *dev, __be32 dst, int scope);
```

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

```
index 09ca529..c270080 100644
```

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

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

```
@@ -784,6 +784,7 @@ int inet_ioctl(struct socket *sock, unsigned int cmd, unsigned long arg)
```

```
{
    struct sock *sk = sock->sk;
    int err = 0;
+ struct net *net = sk->sk_net;
```

```
    switch (cmd) {
        case SIOCGSTAMP:
@@ -795,12 +796,12 @@ int inet_ioctl(struct socket *sock, unsigned int cmd, unsigned long arg)
        case SIOCADDRRT:
        case SIOCDELRT:
        case SIOCRTMSG:
- err = ip_rt_ioctl(sk->sk_net, cmd, (void __user *)arg);
+ err = ip_rt_ioctl(net, cmd, (void __user *)arg);
        break;
```

```

    case SIOCDARP:
    case SIOCGARP:
    case SIOCSARP:
-   err = arp_ioctl(sk->sk_net, cmd, (void __user *)arg);
+   err = arp_ioctl(net, cmd, (void __user *)arg);
    break;
    case SIOCGIFADDR:
    case SIOCSIFADDR:
@@ -813,7 +814,7 @@ int inet_ioctl(struct socket *sock, unsigned int cmd, unsigned long arg)
    case SIOCSIFPFLAGS:
    case SIOCGIFPFLAGS:
    case SIOCSIFFLAGS:
-   err = devinet_ioctl(cmd, (void __user *)arg);
+   err = devinet_ioctl(net, cmd, (void __user *)arg);
    break;
    default:
        if (sk->sk_prot->ioctl)
diff --git a/net/ipv4/devinet.c b/net/ipv4/devinet.c
index 963e711..f7e78b7 100644
--- a/net/ipv4/devinet.c
+++ b/net/ipv4/devinet.c
@@ -595,7 +595,7 @@ static __inline__ int inet_abc_len(__be32 addr)
}

```

```

-int devinet_ioctl(unsigned int cmd, void __user *arg)
+int devinet_ioctl(struct net *net, unsigned int cmd, void __user *arg)
{
    struct ifreq ifr;
    struct sockaddr_in sin_orig;
@@ -624,7 +624,7 @@ int devinet_ioctl(unsigned int cmd, void __user *arg)
    *colon = 0;

#ifdef CONFIG_KMOD
-   dev_load(&init_net, ifr.ifr_name);
+   dev_load(net, ifr.ifr_name);
#endif

    switch (cmd) {
@@ -665,7 +665,7 @@ int devinet_ioctl(unsigned int cmd, void __user *arg)
        rtnl_lock();

        ret = -ENODEV;
-   if ((dev = __dev_get_by_name(&init_net, ifr.ifr_name)) == NULL)
+   if ((dev = __dev_get_by_name(net, ifr.ifr_name)) == NULL)
        goto done;

        if (colon)

```

```
diff --git a/net/ipv4/ipconfig.c b/net/ipv4/ipconfig.c
index a52b585..009d78f 100644
--- a/net/ipv4/ipconfig.c
+++ b/net/ipv4/ipconfig.c
@@ -291,7 +291,7 @@ static int __init ic_dev_ioctl(unsigned int cmd, struct ifreq *arg)

    mm_segment_t oldfs = get_fs();
    set_fs(get_ds());
- res = devinet_ioctl(cmd, (struct ifreq __user *) arg);
+ res = devinet_ioctl(&init_net, cmd, (struct ifreq __user *) arg);
    set_fs(oldfs);
    return res;
}
--
1.5.3.rc5
```
