
Subject: [PATCH 2/5] net: Make rtnetlink infrastructure network namespace aware
Posted by [ebiederm](#) on Sat, 29 Sep 2007 01:02:11 GMT

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

After this patch none of the netlink callback support anything except the initial network namespace but the rtnetlink infrastructure now handles multiple network namespaces.

Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

```
include/linux/rtnetlink.h |  8 +---  
include/net/net_namespace.h |  3 +  
net/bridge/br_netlink.c |  4 +-  
net/core/fib_rules.c |  4 +-  
net/core/neighbour.c |  4 +-  
net/core/rtnetlink.c | 96 ++++++-----  
net/decnet/dn_dev.c |  4 +-  
net/decnet/dn_route.c |  2 +-  
net/decnet/dn_table.c |  4 +-  
net/ipv4/devinet.c |  4 +-  
net/ipv4/fib_semantics.c |  4 +-  
net/ipv4/ipmr.c |  4 +-  
net/ipv4/route.c |  2 +-  
net/ipv6/addrconf.c | 14 +----  
net/ipv6/route.c |  6 +-  
net/sched/act_api.c |  8 +---  
net/sched/cls_api.c |  2 +-  
net/sched/sch_api.c |  4 +-  
net/wireless/wext.c |  5 ++  
19 files changed, 129 insertions(+), 53 deletions(-)
```

```
diff --git a/include/linux/rtnetlink.h b/include/linux/rtnetlink.h  
index 9c21e45..518247e 100644  
--- a/include/linux/rtnetlink.h  
+++ b/include/linux/rtnetlink.h  
@@ -582,11 +582,11 @@ extern int __rtattr_parse_nested_compat(struct rtattr *tb[], int maxattr,  
({ data = RTA_PAYLOAD(rta) >= len ? RTA_DATA(rta) : NULL; \  
  __rtattr_parse_nested_compat(tb, max, rta, len); })  
  
-extern int rtnetlink_send(struct sk_buff *skb, u32 pid, u32 group, int echo);  
-extern int rtnl_unicast(struct sk_buff *skb, u32 pid);  
-extern int rtnl_notify(struct sk_buff *skb, u32 pid, u32 group,  
+extern int rtnetlink_send(struct sk_buff *skb, struct net *net, u32 pid, u32 group, int echo);  
+extern int rtnl_unicast(struct sk_buff *skb, struct net *net, u32 pid);  
+extern int rtnl_notify(struct sk_buff *skb, struct net *net, u32 pid, u32 group,  
    struct nlmsghdr *nlh, gfp_t flags);  
-extern void rtnl_set_sk_err(u32 group, int error);  
+extern void rtnl_set_sk_err(struct net *net, u32 group, int error);
```

```

extern int rtnetlink_put_metrics(struct sk_buff *skb, u32 *metrics);
extern int rtnl_put_cacheinfo(struct sk_buff *skb, struct dst_entry *dst,
    u32 id, u32 ts, u32 tsage, long expires,
diff --git a/include/net/net_namespace.h b/include/net/net_namespace.h
index 934c840..f75607a 100644
--- a/include/net/net_namespace.h
+++ b/include/net/net_namespace.h
@@ -10,6 +10,7 @@

struct proc_dir_entry;
struct net_device;
+struct sock;
struct net {
    atomic_t count; /* To decided when the network
        * namespace should be freed.
@@ -29,6 +30,8 @@ struct net {
    struct list_head dev_base_head;
    struct hlist_head *dev_name_head;
    struct hlist_head *dev_index_head;
+
+ struct sock *rtnl; /* rtinetlink socket */
};

#endif CONFIG_NET
diff --git a/net/bridge/br_netlink.c b/net/bridge/br_netlink.c
index a4ffa2b..f5d6933 100644
--- a/net/bridge/br_netlink.c
+++ b/net/bridge/br_netlink.c
@@ -97,10 +97,10 @@ void br_ifinfo_notify(int event, struct net_bridge_port *port)
    kfree_skb(skb);
    goto errout;
}
- err = rtnl_notify(skb, 0, RTNLGRP_LINK, NULL, GFP_ATOMIC);
+ err = rtnl_notify(skb, &init_net, RTNLGRP_LINK, NULL, GFP_ATOMIC);
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_LINK, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_LINK, err);
}

/*
diff --git a/net/core/fib_rules.c b/net/core/fib_rules.c
index 357bfa0..03c803c 100644
--- a/net/core/fib_rules.c
+++ b/net/core/fib_rules.c
@@ -577,10 +577,10 @@ static void notify_rule_change(int event, struct fib_rule *rule,
    kfree_skb(skb);
    goto errout;
}

```

```

}

- err = rtnl_notify(skb, pid, ops->nlgroup, nlh, GFP_KERNEL);
+ err = rtnl_notify(skb, &init_net, pid, ops->nlgroup, nlh, GFP_KERNEL);
errout:
if (err < 0)
- rtnl_set_sk_err(ops->nlgroup, err);
+ rtnl_set_sk_err(&init_net, ops->nlgroup, err);
}

static void attach_rules(struct list_head *rules, struct net_device *dev)
diff --git a/net/core/neighbour.c b/net/core/neighbour.c
index 27001db..c452584 100644
--- a/net/core/neighbour.c
+++ b/net/core/neighbour.c
@@ @ -2466,10 +2466,10 @@ static void __neigh_notify(struct neighbour *n, int type, int flags)
    kfree_skb(skb);
    goto errout;
}
- err = rtnl_notify(skb, 0, RTNLGRP_NEIGH, NULL, GFP_ATOMIC);
+ err = rtnl_notify(skb, &init_net, 0, RTNLGRP_NEIGH, NULL, GFP_ATOMIC);
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_NEIGH, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_NEIGH, err);
};

#endif CONFIG_ARPD
diff --git a/net/core/rtnetlink.c b/net/core/rtnetlink.c
index 56fc4f9..fc49104 100644
--- a/net/core/rtnetlink.c
+++ b/net/core/rtnetlink.c
@@ @ -60,7 +60,6 @@ struct rtnl_link
};

static DEFINE_MUTEX(rtnl_mutex);
static struct sock *rtnl;

void rtnl_lock(void)
{
@@ @ -74,9 +73,38 @@ void __rtnl_unlock(void)

void rtnl_unlock(void)
{
- mutex_unlock(&rtnl_mutex);
- if (rtnl && rtnl->sk_receive_queue.qlen)
+ struct net *net;
+
+ /*

```

```

+ * Loop through all of the rtnl sockets until none of them (in
+ * a live network namespace) have queue packets.
+ *
+ * We have to be careful with the locking here as
+ * sk_data_ready aka rtneink_rcv takes the rtnl_mutex.
+ *
+ * To ensure the network namespace does not exit while
+ * we are processing packets on it's rtnl socket we
+ * grab a reference to the network namespace, ignoring
+ * it if the network namespace has already exited.
+ */
+retry:
+ for_each_net(net) {
+ struct sock *rtnl = net->rtnl;
+
+ if (!rtnl || !rtnl->sk_receive_queue.qlen)
+ continue;
+
+ if (!maybe_get_net(net))
+ continue;
+
+ mutex_unlock(&rtnl_mutex);
 rtnl->sk_data_ready(rtnl, 0);
+ mutex_lock(&rtnl_mutex);
+ put_net(net);
+ goto retry;
+ }
+ mutex_unlock(&rtnl_mutex);
+
 netdev_run_todo();
}

@@ -462,8 +490,9 @@ size_t rtattr_strlcpy(char *dest, const struct rtattr *rta, size_t size)
 return ret;
}

-int rtneink_send(struct sk_buff *skb, u32 pid, unsigned group, int echo)
+int rtneink_send(struct sk_buff *skb, struct net *net, u32 pid, unsigned group, int echo)
{
+ struct sock *rtnl = net->rtnl;
 int err = 0;

 NETLINK_CB(skb).dst_group = group;
@@ -475,14 +504,17 @@ int rtneink_send(struct sk_buff *skb, u32 pid, unsigned group, int
echo)
 return err;
}

```

```

-int rtnl_unicast(struct sk_buff *skb, u32 pid)
+int rtnl_unicast(struct sk_buff *skb, struct net *net, u32 pid)
{
+ struct sock *rtnl = net->rtnl;
+
 return nlmsg_unicast(rtnl, skb, pid);
}

-int rtnl_notify(struct sk_buff *skb, u32 pid, u32 group,
+int rtnl_notify(struct sk_buff *skb, struct net *net, u32 pid, u32 group,
    struct nlmsghdr *nlh, gfp_t flags)
{
+ struct sock *rtnl = net->rtnl;
    int report = 0;

    if (nlh)
@@ -491,8 +523,10 @@ int rtnl_notify(struct sk_buff *skb, u32 pid, u32 group,
    return nlmsg_notify(rtnl, skb, pid, group, report, flags);
}

-void rtnl_set_sk_err(u32 group, int error)
+void rtnl_set_sk_err(struct net *net, u32 group, int error)
{
+ struct sock *rtnl = net->rtnl;
+
    netlink_set_err(rtnl, 0, group, error);
}

@@ -1205,7 +1239,7 @@ static int rtnl_getlink(struct sk_buff *skb, struct nlmsghdr* nlh, void
*arg)
    kfree_skb(nskb);
    goto errout;
}
- err = rtnl_unicast(nskb, NETLINK_CB(skb).pid);
+ err = rtnl_unicast(nskb, net, NETLINK_CB(skb).pid);
errout:
    dev_put(dev);

@@ -1256,10 +1290,10 @@ void rtmsg_ifinfo(int type, struct net_device *dev, unsigned change)
    kfree_skb(skb);
    goto errout;
}
- err = rtnl_notify(skb, 0, RTNLGRP_LINK, NULL, GFP_KERNEL);
+ err = rtnl_notify(skb, &init_net, 0, RTNLGRP_LINK, NULL, GFP_KERNEL);
errout:
    if (err < 0)
- rtnl_set_sk_err(RTNLGRP_LINK, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_LINK, err);

```

```

}

/* Protected by RTNL sempahore. */
@@ -1270,6 +1304,7 @@ static int rtattr_max;

static int rtnetlink_rcv_msg(struct sk_buff *skb, struct nlmsghdr *nlh)
{
+ struct net *net = skb->sk->sk_net;
    rtnl_doit_func doit;
    int sz_idx, kind;
    int min_len;
@@ -1298,6 +1333,7 @@ static int rtnetlink_rcv_msg(struct sk_buff *skb, struct nlmsghdr *nlh)
    return -EPERM;

    if (kind == 2 && nlh->nlnmsg_flags&NLM_F_DUMP) {
+ struct sock *rtnl;
    rtnl_dumpit_func dumpit;

        dumpit = rtnl_get_dumpit(family, type);
@@ -1305,6 +1341,7 @@ static int rtnetlink_rcv_msg(struct sk_buff *skb, struct nlmsghdr *nlh)
    return -EOPNOTSUPP;

        __rtnl_unlock();
+ rtnl = net->rtnl;
    err = netlink_dump_start(rtnl, skb, nlh, dumpit, NULL);
    rtnl_lock();
    return err;
@@ -1383,6 +1420,40 @@ static struct notifier_block rtnetlink_dev_notifier = {
    .notifier_call = rtnetlink_event,
};

+
+static int rtnetlink_net_init(struct net *net)
+{
+ struct sock *sk;
+ sk = netlink_kernel_create(net, NETLINK_ROUTE, RTNLGRP_MAX,
+ + rtnl_rcv, &rtnl_mutex, THIS_MODULE);
+ if (!sk)
+ return -ENOMEM;
+
+ /* Don't hold an extra reference on the namespace */
+ put_net(sk->sk_net);
+ net->rtnl = sk;
+ return 0;
+}
+
+static void rtnetlink_net_exit(struct net *net)
+{

```

```

+ struct sock *sk = net->rtnl;
+ if (sk) {
+ /* At the last minute lie and say this is a socket for the
+ * initial network namespace. So the socket will be safe to
+ * free.
+ */
+ sk->sk_net = get_net(&init_net);
+ sock_put(sk);
+ net->rtnl = NULL;
+ }
+}
+
+static struct pernet_operations rtinetlink_net_ops = {
+ .init = rtinetlink_net_init,
+ .exit = rtinetlink_net_exit,
+};
+
void __init rtinetlink_init(void)
{
int i;
@@ -1395,10 +1466,9 @@ void __init rtinetlink_init(void)
if (!rt_a_buf)
panic("rtinetlink_init: cannot allocate rt_a_buf\n");

- rtnl = netlink_kernel_create(&init_net, NETLINK_ROUTE, RTNLGRP_MAX,
- rtnl_rcv, &rtnl_mutex, THIS_MODULE);
- if (rtnl == NULL)
+ if (register_pernet_subsys(&rtinetlink_net_ops))
panic("rtinetlink_init: cannot initialize rtinetlink\n");
+
netlink_set_nonroot(NETLINK_ROUTE, NL_NONROOT_RECV);
register_netdevice_notifier(&rtnetlink_dev_notifier);

diff --git a/net/decnet/dn_dev.c b/net/decnet/dn_dev.c
index 6e82f7a..8b34b24 100644
--- a/net/decnet/dn_dev.c
+++ b/net/decnet/dn_dev.c
@@ -791,10 +791,10 @@ static void dn_ifaddr_notify(int event, struct dn_ifaddr *ifa)
kfree_skb(skb);
goto errout;
}
- err = rtnl_notify(skb, 0, RTNLGRP_DECnet_IFADDR, NULL, GFP_KERNEL);
+ err = rtnl_notify(skb, &init_net, 0, RTNLGRP_DECnet_IFADDR, NULL, GFP_KERNEL);
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_DECnet_IFADDR, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_DECnet_IFADDR, err);
}

```

```

static int dn_nl_dump_ifaddr(struct sk_buff *skb, struct netlink_callback *cb)
diff --git a/net/decnet/dn_route.c b/net/decnet/dn_route.c
index a16374b..a5df108 100644
--- a/net/decnet/dn_route.c
+++ b/net/decnet/dn_route.c
@@ -1606,7 +1606,7 @@ static int dn_cache_getroute(struct sk_buff *in_skb, struct nlmsghdr
*nlh, void
    goto out_free;
}
-
```

- return rtnl_unicast(skb, NETLINK_CB(in_skb).pid);

```
+ return rtnl_unicast(skb, &init_net, NETLINK_CB(in_skb).pid);
```



```
out_free:
    kfree_skb(skb);
diff --git a/net/decnet/dn_table.c b/net/decnet/dn_table.c
index a3bdb8d..e09d915 100644
--- a/net/decnet/dn_table.c
+++ b/net/decnet/dn_table.c
@@ -375,10 +375,10 @@ static void dn_rtmsg_fib(int event, struct dn_fib_node *f, int z, u32
tb_id,
    kfree_skb(skb);
    goto errout;
}
-
```

- err = rtnl_notify(skb, pid, RTNLGRP_DECnet_ROUTE, nlh, GFP_KERNEL);

```
+ err = rtnl_notify(skb, &init_net, pid, RTNLGRP_DECnet_ROUTE, nlh, GFP_KERNEL);
```

errout:

```
    if (err < 0)
-    rtnl_set_sk_err(RTNLGRP_DECnet_ROUTE, err);
+    rtnl_set_sk_err(&init_net, RTNLGRP_DECnet_ROUTE, err);
}
```



```
static __inline__ int dn_hash_dump_bucket(struct sk_buff *skb,
diff --git a/net/ipv4/devinet.c b/net/ipv4/devinet.c
index 0ff2ef6..1ae7dcf 100644
--- a/net/ipv4/devinet.c
+++ b/net/ipv4/devinet.c
@@ -1241,10 +1241,10 @@ static void rtmsg_ifa(int event, struct in_ifaddr* ifa, struct nlmsghdr
*nlh,
    kfree_skb(skb);
    goto errout;
}
-
```

- err = rtnl_notify(skb, pid, RTNLGRP_IPV4_IFADDR, nlh, GFP_KERNEL);

```
+ err = rtnl_notify(skb, &init_net, pid, RTNLGRP_IPV4_IFADDR, nlh, GFP_KERNEL);
```

errout:

```
    if (err < 0)
-    rtnl_set_sk_err(RTNLGRP_IPV4_IFADDR, err);
```

```

+ rtnl_set_sk_err(&init_net, RTNLGRP_IPV4_IFADDR, err);
}

#ifndef CONFIG_SYSCTL
diff --git a/net/ipv4/fib_semantics.c b/net/ipv4/fib_semantics.c
index 1351a26..33ec960 100644
--- a/net/ipv4/fib_semantics.c
+++ b/net/ipv4/fib_semantics.c
@@ -320,11 +320,11 @@ void rtmsg_fib(int event, __be32 key, struct fib_alias *fa,
    kfree_skb(skb);
    goto errout;
}
- err = rtnl_notify(skb, info->pid, RTNLGRP_IPV4_ROUTE,
+ err = rtnl_notify(skb, &init_net, info->pid, RTNLGRP_IPV4_ROUTE,
    info->nlh, GFP_KERNEL);
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_IPV4_ROUTE, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_IPV4_ROUTE, err);
}

/* Return the first fib alias matching TOS with
diff --git a/net/ipv4/ipmr.c b/net/ipv4/ipmr.c
index b8b4b49..c077cc0 100644
--- a/net/ipv4/ipmr.c
+++ b/net/ipv4/ipmr.c
@@ -321,7 +321,7 @@ static void ipmr_destroy_unres(struct mfc_cache *c)
    e->error = -ETIMEDOUT;
    memset(&e->msg, 0, sizeof(e->msg));

- rtnl_unicast(skb, NETLINK_CB(skb).pid);
+ rtnl_unicast(skb, &init_net, NETLINK_CB(skb).pid);
} else
    kfree_skb(skb);
}
@@ -533,7 +533,7 @@ static void ipmr_cache_resolve(struct mfc_cache *uc, struct mfc_cache
*c)
    memset(&e->msg, 0, sizeof(e->msg));
}

- rtnl_unicast(skb, NETLINK_CB(skb).pid);
+ rtnl_unicast(skb, &init_net, NETLINK_CB(skb).pid);
} else
    ip_mr_forward(skb, c, 0);
}
diff --git a/net/ipv4/route.c b/net/ipv4/route.c
index a538a30..a710b48 100644
--- a/net/ipv4/route.c

```

```

+++ b/net/ipv4/route.c
@@ -2635,7 +2635,7 @@ static int inet_rtm_getroute(struct sk_buff *in_skb, struct nlmsghdr*
nlh, void
if (err <= 0)
goto errout_free;

- err = rtnl_unicast(skb, NETLINK_CB(in_skb).pid);
+ err = rtnl_unicast(skb, &init_net, NETLINK_CB(in_skb).pid);
errout:
return err;

diff --git a/net/ipv6/addrconf.c b/net/ipv6/addrconf.c
index bdc183e..7dd3f4f 100644
--- a/net/ipv6/addrconf.c
+++ b/net/ipv6/addrconf.c
@@ -3429,7 +3429,7 @@ static int inet6_rtm_getaddr(struct sk_buff *in_skb, struct nlmsghdr*
nlh,
kfree_skb(skb);
goto errout_ifa;
}
- err = rtnl_unicast(skb, NETLINK_CB(in_skb).pid);
+ err = rtnl_unicast(skb, &init_net, NETLINK_CB(in_skb).pid);
errout_ifa:
in6_ifa_put(ifa);
errout:
@@ -3452,10 +3452,10 @@ static void inet6_ifa_notify(int event, struct inet6_ifaddr *ifa)
kfree_skb(skb);
goto errout;
}
- err = rtnl_notify(skb, 0, RTNLGRP_IPV6_IFADDR, NULL, GFP_ATOMIC);
+ err = rtnl_notify(skb, &init_net, 0, RTNLGRP_IPV6_IFADDR, NULL, GFP_ATOMIC);
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_IPV6_IFADDR, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_IPV6_IFADDR, err);
}

static inline void ipv6_store_devconf(struct ipv6_devconf *cnf,
@@ -3660,10 +3660,10 @@ void inet6_ifinfo_notify(int event, struct inet6_dev *idev)
kfree_skb(skb);
goto errout;
}
- err = rtnl_notify(skb, 0, RTNLGRP_IPV6_IFADDR, NULL, GFP_ATOMIC);
+ err = rtnl_notify(skb, &init_net, 0, RTNLGRP_IPV6_IFADDR, NULL, GFP_ATOMIC);
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_IPV6_IFADDR, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_IPV6_IFADDR, err);

```

```

}

static inline size_t inet6_prefix_nlmsg_size(void)
@@ -3729,10 +3729,10 @@ static void inet6_prefix_notify(int event, struct inet6_dev *idev,
    kfree_skb(skb);
    goto errout;
}
- err = rtnl_notify(skb, 0, RTNLGRP_IPV6_PREFIX, NULL, GFP_ATOMIC);
+ err = rtnl_notify(skb, &init_net, 0, RTNLGRP_IPV6_PREFIX, NULL, GFP_ATOMIC);
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_IPV6_PREFIX, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_IPV6_PREFIX, err);
}

static void __ipv6_ifa_notify(int event, struct inet6_ifaddr *ifp)
diff --git a/net/ipv6/route.c b/net/ipv6/route.c
index f0cf11c..849dfd1 100644
--- a/net/ipv6/route.c
+++ b/net/ipv6/route.c
@@ -2307,7 +2307,7 @@ static int inet6_rtm_getroute(struct sk_buff *in_skb, struct nlmsghdr*
nlh, void
    goto errout;
}
- err = rtnl_unicast(skb, NETLINK_CB(in_skb).pid);
+ err = rtnl_unicast(skb, &init_net, NETLINK_CB(in_skb).pid);
errout:
return err;
}
@@ -2337,10 +2337,10 @@ void inet6_rt_notify(int event, struct rt6_info *rt, struct nl_info *info)
    kfree_skb(skb);
    goto errout;
}
- err = rtnl_notify(skb, pid, RTNLGRP_IPV6_ROUTE, nlh, gfp_any());
+ err = rtnl_notify(skb, &init_net, pid, RTNLGRP_IPV6_ROUTE, nlh, gfp_any());
errout:
if (err < 0)
- rtnl_set_sk_err(RTNLGRP_IPV6_ROUTE, err);
+ rtnl_set_sk_err(&init_net, RTNLGRP_IPV6_ROUTE, err);
}

/*
diff --git a/net/sched/act_api.c b/net/sched/act_api.c
index 8528291..8150647 100644
--- a/net/sched/act_api.c
+++ b/net/sched/act_api.c
@@ -660,7 +660,7 @@ act_get_notify(u32 pid, struct nlmsghdr *n, struct tc_action *a, int event)

```

```

    return -EINVAL;
}

- return rtnl_unicast(skb, pid);
+ return rtnl_unicast(skb, &init_net, pid);
}

static struct tc_action *
@@ -781,7 +781,7 @@ static int tca_action_flush(struct rtattr *rta, struct nlmsghdr *n, u32 pid)
nlh->nlmsg_flags |= NLM_F_ROOT;
module_put(a->ops->owner);
kfree(a);
- err = rtnetlink_send(skb, pid, RTNLGRP_TC, n->nlmsg_flags&NLM_F_ECHO);
+ err = rtnetlink_send(skb, &init_net, pid, RTNLGRP_TC, n->nlmsg_flags&NLM_F_ECHO);
if (err > 0)
return 0;

@@ -844,7 +844,7 @@ tca_action_gd(struct rtattr *rta, struct nlmsghdr *n, u32 pid, int event)

/* now do the delete */
tcf_action_destroy(head, 0);
- ret = rtnetlink_send(skb, pid, RTNLGRP_TC,
+ ret = rtnetlink_send(skb, &init_net, pid, RTNLGRP_TC,
          n->nlmsg_flags&NLM_F_ECHO);
if (ret > 0)
return 0;
@@ -888,7 +888,7 @@ static int tcf_add_notify(struct tc_action *a, u32 pid, u32 seq, int event,
nlh->nlmsg_len = skb_tail_pointer(skb) - b;
NETLINK_CB(skb).dst_group = RTNLGRP_TC;

- err = rtnetlink_send(skb, pid, RTNLGRP_TC, flags&NLM_F_ECHO);
+ err = rtnetlink_send(skb, &init_net, pid, RTNLGRP_TC, flags&NLM_F_ECHO);
if (err > 0)
err = 0;
return err;
diff --git a/net/sched/cls_api.c b/net/sched/cls_api.c
index 2754e50..9dc88 100644
--- a/net/sched/cls_api.c
+++ b/net/sched/cls_api.c
@@ -361,7 +361,7 @@ static int tfilter_notify(struct sk_buff *oskb, struct nlmsghdr *n,
    return -EINVAL;
}

- return rtnetlink_send(skb, pid, RTNLGRP_TC, n->nlmsg_flags&NLM_F_ECHO);
+ return rtnetlink_send(skb, &init_net, pid, RTNLGRP_TC, n->nlmsg_flags&NLM_F_ECHO);
}

struct tcf_dump_args

```

```

diff --git a/net/sched/sch_api.c b/net/sched/sch_api.c
index 95eb0a8..ddbae5a 100644
--- a/net/sched/sch_api.c
+++ b/net/sched/sch_api.c
@@ -872,7 +872,7 @@ static int qdisc_notify(struct sk_buff *oskb, struct nlmsghdr *n,
}

if (skb->len)
- return rtnetlink_send(skb, pid, RTNLGRP_TC, n->nmsg_flags&NLM_F_ECHO);
+ return rtnetlink_send(skb, &init_net, pid, RTNLGRP_TC, n->nmsg_flags&NLM_F_ECHO);

err_out:
kfree_skb(skb);
@@ -1103,7 +1103,7 @@ static int tclass_notify(struct sk_buff *oskb, struct nlmsghdr *n,
    return -EINVAL;
}

- return rtnetlink_send(skb, pid, RTNLGRP_TC, n->nmsg_flags&NLM_F_ECHO);
+ return rtnetlink_send(skb, &init_net, pid, RTNLGRP_TC, n->nmsg_flags&NLM_F_ECHO);
}

struct qdisc_dump_args
diff --git a/net/wireless/wext.c b/net/wireless/wext.c
index 85e5f9d..85805f2 100644
--- a/net/wireless/wext.c
+++ b/net/wireless/wext.c
@@ -1137,7 +1137,7 @@ static void wireless_nlevent_process(unsigned long data)
    struct sk_buff *skb;

    while ((skb = skb_dequeue(&wireless_nlevent_queue)))
- rtnl_notify(skb, 0, RTNLGRP_LINK, NULL, GFP_ATOMIC);
+ rtnl_notify(skb, &init_net, 0, RTNLGRP_LINK, NULL, GFP_ATOMIC);
}

static DECLARE_TASKLET(wireless_nlevent_tasklet, wireless_nlevent_process, 0);
@@ -1189,6 +1189,9 @@ static void rtmmsg_iwinfo(struct net_device *dev, char *event, int
event_len)
    struct sk_buff *skb;
    int err;

+ if (dev->nd_net != &init_net)
+ return;
+
    skb = nlmsg_new(NLMSG_DEFAULT_SIZE, GFP_ATOMIC);
    if (!skb)
        return;
--
```

1.5.3.rc6.17.g1911

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
