
Subject: [patch 04/12] net namespace : isolate the inet device.
Posted by [Daniel Lezcano](#) on Fri, 19 Jan 2007 15:47:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Daniel Lezcano <dlezcano@fr.ibm.com>

ip and ifconfig commands will not show ip addr
not belonging to the current network namespace.

Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>

include/linux/inetdevice.h | 1 +
net/ipv4/devinet.c | 22 ++++++
2 files changed, 22 insertions(+), 1 deletion(-)

Index: 2.6.20-rc4-mm1/include/linux/inetdevice.h

```
=====
--- 2.6.20-rc4-mm1.orig/include/linux/inetdevice.h
+++ 2.6.20-rc4-mm1/include/linux/inetdevice.h
@@ -99,6 +99,7 @@
     unsigned char ifa_flags;
     unsigned char ifa_prefixlen;
     char ifa_label[IFNAMSIZ];
+ struct net_namespace *ifa_net_ns;
};
```

```
extern int register_inetaddr_notifier(struct notifier_block *nb);
```

Index: 2.6.20-rc4-mm1/net/ipv4/devinet.c

```
=====
--- 2.6.20-rc4-mm1.orig/net/ipv4/devinet.c
+++ 2.6.20-rc4-mm1/net/ipv4/devinet.c
@@ -53,6 +53,7 @@
#include <linux/notifier.h>
#include <linux/inetdevice.h>
#include <linux/igmp.h>
+#include <linux/net_namespace.h>
#ifdef CONFIG_SYSCTL
#include <linux/sysctl.h>
#endif
@@ -269,6 +270,7 @@

if (!(ifa->ifa_flags & IFA_F_SECONDARY) ||
    ifa1->ifa_mask != ifa->ifa_mask ||
+ !net_ns_match(ifa->ifa_net_ns, ifa1->ifa_net_ns) ||
    !inet_ifa_match(ifa1->ifa_address, ifa)) {
    ifap1 = &ifa->ifa_next;
    prev_prom = ifa;
```

@@ -471,6 +473,9 @@

```
for (ifap = &in_dev->ifa_list; (ifa = *ifap) != NULL;
    ifap = &ifa->ifa_next) {
+ if (!net_ns_match(ifa->ifa_net_ns, current_net_ns))
+ continue;
+
    if (tb[IFA_LOCAL] &&
        ifa->ifa_local != nla_get_be32(tb[IFA_LOCAL]))
        continue;
```

@@ -544,6 +549,7 @@

```
    ifa->ifa_flags = ifm->ifa_flags;
    ifa->ifa_scope = ifm->ifa_scope;
    ifa->ifa_dev = in_dev;
+ ifa->ifa_net_ns = current_net_ns;
```

```
    ifa->ifa_local = nla_get_be32(tb[IFA_LOCAL]);
    ifa->ifa_address = nla_get_be32(tb[IFA_ADDRESS]);
```

@@ -689,6 +695,8 @@

```
    for (ifap = &in_dev->ifa_list; (ifa = *ifap) != NULL;
        ifap = &ifa->ifa_next) {
        if (!strcmp(ifr.ifr_name, ifa->ifa_label) &&
+         net_ns_match(ifa->ifa_net_ns,
+         current_net_ns) &&
            sin_orig.sin_addr.s_addr ==
            ifa->ifa_address) {
            break; /* found */
```

@@ -701,11 +709,16 @@

```
    if (!ifa) {
        for (ifap = &in_dev->ifa_list; (ifa = *ifap) != NULL;
            ifap = &ifa->ifa_next)
-         if (!strcmp(ifr.ifr_name, ifa->ifa_label))
+         if (!strcmp(ifr.ifr_name, ifa->ifa_label) &&
+             net_ns_match(ifa->ifa_net_ns,
+             current_net_ns))
            break;
    }
}
```

```
+ if (ifa && !net_ns_match(ifa->ifa_net_ns, current_net_ns))
```

```
+ goto done;
```

```
+
```

```
ret = -EADDRNOTAVAIL;
if (!ifa && cmd != SIOCSIFADDR && cmd != SIOCSIFFLAGS)
goto done;
```

@@ -749,6 +762,8 @@

```
ret = -ENOBUFS;
if ((ifa = inet_alloc_ifa()) == NULL)
```

```

    break;
+
+ ifa->ifa_net_ns = current_net_ns;
  if (colon)
    memcpy(ifa->ifa_label, ifr.ifr_name, IFNAMSIZ);
  else
@@ -853,6 +868,8 @@
    goto out;

    for (; ifa; ifa = ifa->ifa_next) {
+ if (!net_ns_match(ifa->ifa_net_ns, current_net_ns))
+ continue;
    if (!buf) {
      done += sizeof(ifr);
      continue;
@@ -1086,6 +1103,7 @@
      in_dev_hold(in_dev);
      ifa->ifa_dev = in_dev;
      ifa->ifa_scope = RT_SCOPE_HOST;
+ ifa->ifa_net_ns = current_net_ns;
      memcpy(ifa->ifa_label, dev->name, IFNAMSIZ);
      inet_insert_ifa(ifa);
    }
@@ -1198,6 +1216,8 @@

    for (ifa = in_dev->ifa_list, ip_idx = 0; ifa;
         ifa = ifa->ifa_next, ip_idx++) {
+ if (!net_ns_match(ifa->ifa_net_ns, current_net_ns))
+ continue;
    if (ip_idx < s_ip_idx)
      continue;
    if (inet_fill_ifaddr(skb, ifa, NETLINK_CB(cb->skb).pid,

```

--

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