
Subject: [patch 0/2] Network namespace L3 : broadcast and fix
Posted by [Daniel Lezcano](#) on Fri, 15 Dec 2006 17:20:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

They apply to 2.6.19-rc6-mm2 as well as 2.6.19-mm1

--

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

Subject: [patch 1/2] net namespace : Add broadcasting
Posted by [Daniel Lezcano](#) on Fri, 15 Dec 2006 17:20:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

Broadcast packets should be delivered to l2 and all l3 childs

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

```
include/linux/net_namespace.h | 10 ++++++++
net/core/net_namespace.c    | 16 ++++++-----
net/ipv4/udp.c            |  5 +++
3 files changed, 29 insertions(+), 2 deletions(-)
```

Index: 2.6.19-mm1/include/linux/net_namespace.h

```
--- 2.6.19-mm1.orig/include/linux/net_namespace.h
```

```
+++ 2.6.19-mm1/include/linux/net_namespace.h
```

```
@@ -9,6 +9,7 @@
```

```
struct in_ifaddr;
struct sk_buff;
+struct sock;
```

```
struct net_ns_net_device {
    struct hlist_head *name_head;
@@ -103,6 +104,9 @@ extern struct net_namespace
```

```
extern int net_ns_ifa_is_visible(const struct in_ifaddr *ifa);
```

```
+extern int net_ns_sock_is_visible(const struct sock *sk,
+    const struct net_namespace *net_ns);
+
```

```

#define SELECT_SRC_ADDR net_ns_select_source_address

#else /* CONFIG_NET_NS */
@@ -178,6 +182,12 @@ static inline int net_ns_ifa_is_visible(
    return 1;
}

+static inline int net_ns_sock_is_visible(const struct sock *sk,
+    const struct net_namespace *net_ns)
+{
+    return 1;
+}
+
#define SELECT_SRC_ADDR inet_select_addr

#endif /* !CONFIG_NET_NS */
Index: 2.6.19-mm1/net/core/net_namespace.c
=====
--- 2.6.19-mm1.orig/net/core/net_namespace.c
+++ 2.6.19-mm1/net/core/net_namespace.c
@@ -11,12 +11,14 @@
#include <linux/net_namespace.h>
#include <linux/net.h>
#include <linux/netdevice.h>
-#include <net/ip_fib.h>
#include <linux/inetdevice.h>
#include <linux/in.h>
#include <linux/skbuff.h>
#include <linux/ip.h>

+#include <net/ip_fib.h>
+#include <net/sock.h>
+
static spinlock_t net_ns_list_lock = SPIN_LOCK_UNLOCKED;

struct net_namespace init_net_ns = {
@@ -424,4 +426,16 @@ int net_ns_ifa_is_visible(const struct i
    return 0;
}

+int net_ns_sock_is_visible(const struct sock *sk,
+    const struct net_namespace *net_ns)
+{
+    if (net_ns->level == NET_NS_LEVEL3)
+        net_ns = net_ns->parent;
+
+    if (sk->sk_net_ns->level == NET_NS_LEVEL3)
+        return sk->sk_net_ns->parent == net_ns;

```

```

+ else
+   return sk->sk_net_ns == net_ns;
+
+}
#endif /* CONFIG_NET_NS */
Index: 2.6.19-mm1/net/ipv4/udp.c
=====
--- 2.6.19-mm1.orig/net/ipv4/udp.c
+++ 2.6.19-mm1/net/ipv4/udp.c
@@ -306,9 +306,12 @@ static inline struct sock *udp_v4_mcast_
    (inet->dport != rmt_port && inet->dport) ||
    (inet->rcv_saddr && inet->rcv_saddr != loc_addr) ||
    ipv6_only_sock(s) ||
-   !net_ns_match(sk->sk_net_ns, ns) ||
-   (s->sk_bound_dev_if && s->sk_bound_dev_if != dif))
    continue;
+#ifdef CONFIG_NET_NS
+  if (!net_ns_sock_is_visible(sk, ns))
+    continue;
+#endif
  if (!ip_mc_sf_allow(s, loc_addr, rmt_addr, dif))
    continue;
  goto found;

--
```

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

Subject: [patch 2/2] net namespace : fix bad hash dev list initialization
 Posted by [Daniel Lezcano](#) on Fri, 15 Dec 2006 17:20:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

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

net/core/dev.c | 5 +---
 1 files changed, 3 insertions(+), 2 deletions(-)

Index: 2.6.19-mm1/net/core/dev.c
=====

--- 2.6.19-mm1.orig/net/core/dev.c
+++ 2.6.19-mm1/net/core/dev.c
@@ -188,6 +188,7 @@ EXPORT_SYMBOL(dev_base);

```
DEFINE_RWLOCK(dev_base_lock);
EXPORT_SYMBOL(dev_base_lock);

+#define dev_hash_size (sizeof(struct hlist_head)*(1<<NETDEV_HASHBITS))
#ifndef CONFIG_NET_NS
#define dev_name_head (current_net_ns->net_device.name_head)
#define dev_index_head (current_net_ns->net_device.index_head)
@@ -3647,10 +3648,10 @@ static int __init net_dev_init(void)
    for (i = 0; i < 16; i++)
        INIT_LIST_HEAD(&ptype_base[i]);

- for (i = 0; i < ARRAY_SIZE(dev_name_head); i++)
+ for (i = 0; i < dev_hash_size; i++)
    INIT_HLIST_HEAD(&dev_name_head[i]);

- for (i = 0; i < ARRAY_SIZE(dev_index_head); i++)
+ for (i = 0; i < dev_hash_size; i++)
    INIT_HLIST_HEAD(&dev_index_head[i]);

/*
--
```

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