

Subject: [PATCH 11/11 net-2.6.26] [NETNS]: Use TCP control socket from a correct namespace.

Posted by [den](#) on Fri, 28 Mar 2008 09:08:58 GMT

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

---

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

---

net/ipv4/tcp\_ipv4.c | 10 ++++-----

1 files changed, 4 insertions(+), 6 deletions(-)

diff --git a/net/ipv4/tcp\_ipv4.c b/net/ipv4/tcp\_ipv4.c

index 14944af..34f9174 100644

--- a/net/ipv4/tcp\_ipv4.c

+++ b/net/ipv4/tcp\_ipv4.c

@@ -88,9 +88,6 @@ int sysctl\_tcp\_low\_latency \_\_read\_mostly;

/\* Check TCP sequence numbers in ICMP packets. \*/

#define ICMP\_MIN\_LENGTH 8

/\* Socket used for sending RSTs \*/

-static struct sock \*tcp\_sock \_\_read\_mostly;

-

void tcp\_v4\_send\_check(struct sock \*sk, int len, struct sk\_buff \*skb);

#ifdef CONFIG\_TCP\_MD5SIG

@@ -598,7 +595,8 @@ static void tcp\_v4\_send\_reset(struct sock \*sk, struct sk\_buff \*skb)  
sizeof(struct tcphdr), IPPROTO\_TCP, 0);

arg.csumoffset = offsetof(struct tcphdr, check) / 2;

- ip\_send\_reply(tcp\_sock, skb, &arg, arg.iov[0].iov\_len);

+ ip\_send\_reply(dev\_net(skb->dst->dev)->ipv4.tcp\_sock, skb,

+ &arg, arg.iov[0].iov\_len);

TCP\_INC\_STATS\_BH(TCP\_MIB\_OUTSEGS);

TCP\_INC\_STATS\_BH(TCP\_MIB\_OUTRSTS);

@@ -693,7 +691,8 @@ static void tcp\_v4\_send\_ack(struct tcp\_timewait\_sock \*twsk,  
if (twsk)

arg.bound\_dev\_if = twsk->tw\_sk.tw\_bound\_dev\_if;

- ip\_send\_reply(tcp\_sock, skb, &arg, arg.iov[0].iov\_len);

+ ip\_send\_reply(dev\_net(skb->dev)->ipv4.tcp\_sock, skb,

+ &arg, arg.iov[0].iov\_len);

TCP\_INC\_STATS\_BH(TCP\_MIB\_OUTSEGS);

}

@@ -2517,7 +2516,6 @@ void \_\_init tcp\_v4\_init(void)

{

if (register\_pernet\_device(&tcp\_sk\_ops))

panic("Failed to create the TCP control socket.\n");

```
- tcp_sock = init_net.ipv4.tcp_sock;  
}
```

```
EXPORT_SYMBOL(ipv4_specific);
```

```
--
```

1.5.3.rc5

---

Containers mailing list

Containers@lists.linux-foundation.org

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

---