

---

Subject: Re: [RFC] L3 network isolation : broadcast  
Posted by [Daniel Lezcano](#) on Fri, 15 Dec 2006 10:04:14 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Daniel Lezcano wrote:

> Hi all,  
>  
> I am trying to find a solution to handle the broadcast traffic on the I3  
> namespace.  
>  
> The broadcast issue comes from the I2 isolation:  
>  
> in udp.c  
>  
> static inline struct sock \*udp\_v4\_mcast\_next(struct sock \*sk,  
>     \_\_be16 loc\_port,  
>     \_\_be32 loc\_addr,  
>     \_\_be16 rmt\_port,  
>     \_\_be32 rmt\_addr,  
>     int dif)  
> {  
>     struct hlist\_node \*node;  
>     struct sock \*s = sk;  
>     struct net\_namespace \*ns = current\_net\_ns;  
>     unsigned short hnum = ntohs(loc\_port);  
>  
>     sk\_for\_each\_from(s, node) {  
>         struct inet\_sock \*inet = inet\_sk(s);  
>  
>         if (inet->num != hnum   ||  
>             (inet->daddr && inet->daddr != rmt\_addr) ||  
>             (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;  
>         if (!ip\_mc\_sf\_allow(s, loc\_addr, rmt\_addr, dif))  
>         continue;  
>         goto found;  
>     }  
>     s = NULL;  
> found:  
>     return s;  
> }  
>  
> This is absolutely correct for I2 namespaces because they share the  
> socket hash table. But that is not correct for I3 namespaces because we

```
> want to deliver the packet to each l3 namespaces which have binded to  
> the broadcast address, so we should avoid checking net_ns_match if we  
> are in a layer 3 namespace. Doing that we will break the l2 isolation  
> because an another l2 namespace could have binded to the same broadcast  
> address.  
>  
> The solution I see here is:  
>  
> if namespace is l3 then;  
>   net_ns match any net_ns registered as listening on this address  
> else  
>   net_ns_match  
> fi
```

Finally, I found a more simple solution which does not need extra registering address and is more efficient.

The packet is to be delivered to the l2 namespace and all l3 childs.  
So, instead of doing net\_ns\_match in the udp\_v4\_mcast\_next, we do net\_ns\_sock\_is\_visible and this function does:

```
if current_namespace is l3 then;  
namespace = current_namespace->parent  
fi  
  
if socket->namespace is l3 then  
return socket->namespace->parent == namespace  
else  
return socket->namespace == namespace  
fi
```

---