
Subject: Re: [PATCH 1/4] netns: Tag the network flow with the network namespace it is in (v2)

Posted by [davem](#) on Wed, 05 Dec 2007 10:13:18 GMT

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

From: "Denis V. Lunev" <den@openvz.org>

Date: Tue, 4 Dec 2007 12:53:33 +0300

> As well as marking flows this indirectly marks the ipv4 routing cache
> as every routing entry contains a flow.
>
> It is useful to add the network namespace into flows as frequently
> the routing information for ingoing and outgoing network packets is
> collected into a flow structure which is then used for several functions
> as it sorts out what is going on.
>
> Changes from v1:
> - remove flow.h dependency from net_namespace.h
>
> Signed-off-by: Denis V. Lunev <den@openvz.org>
> Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

Hmmm, actually I change my mind.

> @@ -10,7 +10,9 @@
> #include <linux/in6.h>
> #include <asm/atomic.h>
>
> +struct net;
> struct flowi {
> + struct net *fl_net;
> int oif;
> int iif;
> __u32 mark;

I'm not applying this, it's going to have a negative impact on routing performance.

It also changes the semantics of the flowi object in a way I very much dislike, in that there is now non-clobberable state in there.

Previously only addressing identifying objects were present in the flow, you could use it any context, and there were no pointer dereferencing or object references from this thing. It was very simple.

That is no longer the case after your patch and I don't want us to go down this path.

Please find another way to implement this.
