
Subject: Re: [PATCH] iptables 32bit compat layer
Posted by [Mishin Dmitry](#) on Wed, 29 Mar 2006 12:38:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wednesday 29 March 2006 16:32, Patrick McHardy wrote:

```
> Dmitry Mishin wrote:
> > On Wednesday 29 March 2006 13:28, Patrick McHardy wrote:
> >>>diff --git a/net/compat.c b/net/compat.c
> >>>index 13177a1..6a7028e 100644
> >>>--- a/net/compat.c
> >>>+++ b/net/compat.c
> >>>@@ -476,8 +476,7 @@ asmlinkage long compat_sys_setsockopt(in
> >>> int err;
> >>> struct socket *sock;
> >>>
> >>>- /* SO_SET_REPLACE seems to be the same in all levels */
> >>>- if (optname == IPT_SO_SET_REPLACE)
> >>>+ if (level == SOL_IPV6 && optname == IPT_SO_SET_REPLACE)
> >>> return do_netfilter_replace(fd, level, optname,
> >>>     optval, optlen);
> >>
> >>I don't understand the reason for this change. If its not a mistake,
> >>it would make more sense to check for IP6T_SO_SET_REPLACE I guess ..
> >
> > IP6T_SO_SET_REPLACE == IPT_SO_SET_REPLACE == XT_SO_SET_REPLACE.
> > Rename will require respective #include directive rename, so, I just
> > leave this as it is. BTW, I'll make respective patch for IPV6 in the near
> > future and this hunk will be removed at all.
>
> I know, but SOL_IPV6 implies IP6T_* - but please don't bother sending
> a new patch for this :) So the point of the change is to exclude IPv6
> from the compat layer because its not implemented yet?
Exactly. Because do_netfilter_replace still works for some cases, but newer
replacement isn't ready yet.
```

--
Thanks,
Dmitry.
