

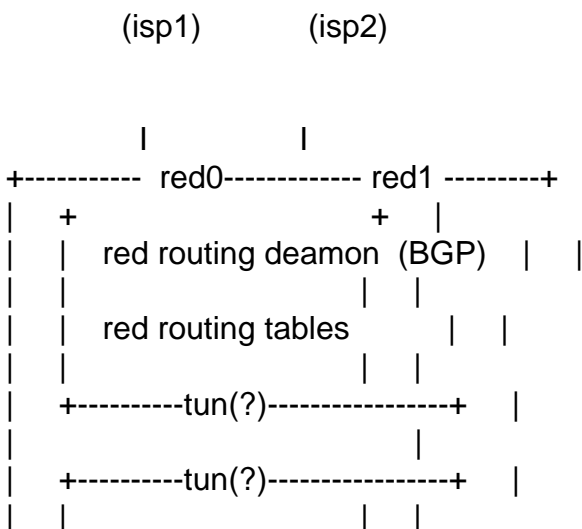
Subject: Routing tables (Re: [patch 2/6] [Network namespace] Network device sharing by view)


Posted by [Kari Hurtta](#) on Thu, 06 Jul 2006 09:45:29 GMT

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

> Andrey Savochkin wrote:
> > Daniel,
> >
> > On Mon, Jun 26, 2006 at 05:49:41PM +0200, Daniel Lezcano wrote:
> >
> >>>Then you lose the ability for each namespace to have its own routing entries.
> >>>Which implies that you'll have difficulties with devices that should exist
> >>>and be visible in one namespace only (like tunnels), as they require IP
> >>>addresses and route.
> >>
> >>I mean instead of having the route tables private to the namespace, the
> >>routes have the information to which namespace they are associated.
> >
> >
> > I think I understand what you're talking about: you want to make routing
> > responsible for determining destination namespace ID in addition to route
> > type (local, unicast etc), nexthop information, and so on. Right?
>
> Yes.
>
> >
> > My point is that if you make namespace tagging at routing time, and
> > your packets are being routed only once, you lose the ability
> > to have separate routing tables in each namespace.
>
> Right. What is the advantage of having separate the routing tables ?

One application may be following. Consider firewall





That may allow running different routing deamon on red and green side. That is possible if they manage different routing tables on kernel. They not need communigate together, when route between them is static.

/ Kari Hurtta

- > -
- > To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
- > the body of a message to majordomo@vger.kernel.org
- > More majordomo info at <http://vger.kernel.org/majordomo-info.html>
- > Please read the FAQ at <http://www.tux.org/lkml/>