
Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by [Andrey Savochkin](#) on Tue, 27 Jun 2006 16:19:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

Herbert,

On Tue, Jun 27, 2006 at 05:48:19PM +0200, Herbert Poetzl wrote:

> On Tue, Jun 27, 2006 at 01:09:11PM +0400, Andrey Savochkin wrote:

> >

> > On Mon, Jun 26, 2006 at 10:02:25PM +0200, Herbert Poetzl wrote:

> > >

> > > - traffic between guests

> > > can be as high (or even higher) than the

> > > outbound traffic, just think web guest and

> > > database guest

> >

> > My experience with host-guest systems tells me the opposite: outside

> > traffic is a way higher than traffic between guests. People put web

> > server and database in different guests not more frequent than they

> > put them on separate physical server. Unless people are building a

> > really huge system when 1 server can't take the whole load, web and

> > database live together and benefit from communications over UNIX

> > sockets.

>

> well, that's probably because you (or your company)

> focuses on providers which simply (re)sell the entities

> to their customers, in which case it would be more

> expensive to put e.g. the database into a separate

> guest. but let me state here that this is not the only

> application for this technology

I'm just sharing my experience.

You have one experience, I have another, and your classification of traffic importance is not the universal one.

My point was that we shouldn't overestimate the use of INET sockets vs. UNIX ones in configurations where communications but not web/db operations play a big role in overall performance.

And indeed I've talked with many different people, from universities to large enterprises.

>

[snip]

> > I'd like to caution about over-optimizing communications between

> > different network namespaces. Many optimizations of local traffic

> > (such as high MTU) don't look so appealing when you start to think

> > about live migration of namespaces.

>

> I think the 'optimization' (or to be precise: desire

- > not to sacrifice local/loopback traffic for some use
- > case as you describe it) does not interfere with live
- > migration at all, we still will have 'local' and 'remote'
- > traffic, and personally I doubt that the live migration
- > is a feature for the masses ...

Why not for the masses?

Andrey
