

Herbert,

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

- >
- > keep in mind that you actually have three kinds
- > of network traffic on a typical host/guest system:
- >
- > - traffic between unit and outside
- > - host traffic should be quite minimal
- > - guest traffic will be quite high
- >
- > - traffic between host and guest
- > probably minimal too (only for shared services)
- >
- > - 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.

Guests are usually comprised of web-db pairs, and people place many such
guests on a single computer.

- >
- > > The routing between network namespaces does have the potential to be
- > > more expensive than just a packet trivially coming off the wire into a
- > > socket.
- >
- > IMHO the routing between network namespaces should
- > not require more than the current local traffic
- > does (i.e. you should be able to achieve loopback
- > speed within an insignificant tolerance) and not
- > nearly the time required for on-wire stuff ...

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.

Regards
Andrey
