

---

Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view  
Posted by [ebiederm](#) on Tue, 27 Jun 2006 16:40:18 GMT

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

---

Herbert Poetzl <herbert@13thfloor.at> writes:

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

>>

>> 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 ...

Several things.

- The linux loopback device is not strongly optimized, it is a compatibility layer.
- Traffic between guests is an implementation detail.  
There is nothing fundamental in our semantics that says the traffic has to be slow for any workload (except for the limits imposed by using actual on the wire protocols). The lo shares the same problem.

Worry about this case now when it has clearly been shown that there are several possible ways to optimize this and get back any lost local performance is optimizing way too early.

Criticize the per namespace performance and all you want. That is pretty much a merge blocker. Unless we do worse than a 1-5% penalty the communication across namespaces is really a non-issue.

Even with your large communications flows between guests 1-5% is nothing.

Eric

---