

jamal <hadi@cyberus.ca> writes:

> Andrey,  
>  
> On Wed, 2006-28-06 at 18:19 +0400, Andrey Savochkin wrote:  
>> Hi Jamal,  
>>  
>> On Wed, Jun 28, 2006 at 09:53:23AM -0400, jamal wrote:  
>> >  
>  
>>  
>> Seeing guestXX-eth0 interfaces by standard tools has certain attractive  
>> sides. But it creates a lot of undesired side effects.  
>>  
>  
> I apologize because i butted into the discussion without perhaps reading  
> the full thread.

This thread is serving as an educational vehicle, and the more people from outside of our little biased group that begin to understand what we are about the better.

>> For example, ntpd queries all network devices by the same ioctls as ifconfig,  
>> and creates separate sockets bound to IP addresses of each device, which is  
>> certainly not desired with namespaces.  
>>  
>  
> Ok, so the problem is that ntp in this case runs on the host side as  
> opposed to the guest? This would explain why Eric is reacting vehemently  
> to the suggestion.

Yes, that would be one problem case.

>> Or more subtle question: do you want hotplug events to be generated when  
>> guest0-eth0 interface comes up in the root namespace, and standard scripts  
>> to try to set some IP address on this interface?..  
>>  
>  
> yes, thats what i was thinking. Even go further and actually create  
> guestxx-eth0 on the host (which results in creating eth0 on the guest)  
> and other things.  
>  
>> In my opinion, the downside of this scheme overweights possible advantages,  
>> and I'm personally quite happy with running commands with switched namespace,

>> like  
>> vzctl exec guest0 ip addr list  
>> vzctl exec guest0 ip link set eth0 up  
>> and so on.  
>  
> Ok, above may be good enough and doesnt require any state it seems on  
> the host side.  
> I got motivated when the word "migration" was mentioned. I understood it  
> to be meaning that a guest may become inoperative for some reason and  
> that its info will be transfered to another guest which may be local or  
> even remote. In such a case, clearly one would need a protocol and the  
> state of all guests sitting at the host. Maybe i am over-reaching.

Not really. Network namespaces while useful in their own right lay the foundation for some more interesting applications. Application migration between machines in particular.

The biggest fundamental problem in migration is after checkpointing your application you can not acquire the resources you need on the new machine because of name conflicts.

So for those of us concerned with migration a question we ask is can we successfully import resource names that another machine has assigned without consulting us.

The context for all of this goes to other discussion that we have been having since January. Breaking all of this into small pieces that can be merged and tested a little at a time is a challenge.

Eric

---