
Subject: Re: [RFC] network namespaces

Posted by [Cedric Le Goater](#) on Wed, 06 Sep 2006 20:25:10 GMT

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

Eric W. Biederman wrote:

>> This family of containers are used too for HPC (high performance computing) and
>> for distributed checkpoint/restart. The cluster runs hundred of jobs, spawning
>> them on different hosts inside an application container. Usually the jobs
>> communicates with broadcast and multicast.
>> Application containers does not care of having different MAC address and rely on
>> a layer 3 approach.

>

> Ok I think to understand this we need some precise definitions.

> In the normal case it is an error for a job to communication with a different
> job.

hmm ? What about an MPI application ?

I would expect each MPI task to be run in its container on different nodes
or on the same node. These individual tasks communicate between each
other through the MPI layer (not only TCP btw) to complete a large calculation.

> The basic advantage with a different MAC is that you can found out who the
> intended recipient is sooner in the networking stack and you have truly
> separate network devices. Allowing for a cleaner implementation.

>

> Changing the MAC after migration is likely to be fine.

indeed.

C.
