Subject: My netns patches updated to Linus' latest
Posted by ebiederm on Thu, 19 Jul 2007 07:05:32 GMT
View Forum Message <> Reply to Message

Ok. Now that I finally have a working version of sysfs I can
start worrying about what else needs to happen to start merging
the network namespace work.

Currently I have two interesting branches available (all essentially
with the same content)

Against the latest -mm
git://git.kernel.org/pub/scm/linux/kernel/git/ebiederm/linux-2.6-netns.git#netns/v2.6.22-rc6-mm1-netns23

Against a snapshot of Linus' current tree.
git://git.kernel.org/pub/scm/linux/kernel/git/ebiederm/linux-2.6-netns.git#netns/v2.6.22netns24

The code has been reduced to a simple struct net, and a few other
tweaks like removing dev_net_lock from the network namespace structure
have been performed.

There is a config option but all it does at present is enable/disable
the option of cloning the network namespace.

My next project is most likely going to be to dig in and make the
loopback device a normal network device (not statically allocated)
as a cleanup in preparation for the network namespace work.

Feel free to pick the patches apart. Especially in the core of the
network stack.

I'm hoping to start seriously submitting patches short after the merge
window closes. So we can get a long test/review cycle.

Eric

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers

Subject: Re: My netns patches updated to Linus' latest
Posted by den on Thu, 19 Jul 2007 13:25:39 GMT
View Forum Message <> Reply to Message

Eric W. Biederman wrote:
> My next project is most likely going to be to dig in and make the

> loopback device a normal network device (not statically allocated)
> as a cleanup in preparation for the network namespace work.

I can do this :)

As far as I understand the code, loopback device is used in two fashions:
- as a real network device
- and as a "stable" never freed piece of memory
I presume, they should be separated.

As for device itself, there is a some sort of a template in a Pavel's
virtual ethernet device driver...

> Feel free to pick the patches apart. Especially in the core of the
> network stack.

I think that a small plan of submitting process is a good idea, at least
we will be able to coordinate the efforts and work in parallel.

> I'm hoping to start seriously submitting patches short after the merge
> window closes. So we can get a long test/review cycle.
:)

Regards,
 Den

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers

---

Subject: Re: My netns patches updated to Linus' latest
Posted by ebiederm on Thu, 19 Jul 2007 13:53:16 GMT
View Forum Message <> Reply to Message

"Denis V. Lunev" <den@sw.ru> writes:

> Eric W. Biederman wrote:
>> My next project is most likely going to be to dig in and make the
>> loopback device a normal network device (not statically allocated)
>> as a cleanup in preparation for the network namespace work.
>
> I can do this :)
>

> As far as I understand the code, loopback device is used in two fashions:
> - as a real network device
> - and as a "stable" never freed piece of memory
> I presume, they should be separated.

Essentially. As far as that goes I think my patchset already covers that aspect of it. Although there may be an initialization or shutdown races that I currently do not handle.

Right now I'm working on a building up a clean patchset, and one of the basic rules is that cleanups should happen before features are added. Which means that before we move a pointer to the loopback device into a per network namespace structure we should make it dynamically allocated if possible.

> As for device itself, there is a some sort of a template in a Pavel's
> virtual ethernet device driver...

Yes. The tricky bits are mostly in the code review for making certain that the loopback device isn't used after it is freed.

>> Feel free to pick the patches apart. Especially in the core of the
>> network stack.
>
> I think that a small plan of submitting process is a good idea, at least
> we will be able to coordinate the efforts and work in parallel.

Yes. Working in parallel would be good. However the bottleneck I see is is patch submission, review and acceptance. Not so much development of working patches.

So as much internal review and testing we can do before the patches get to Dave Miller the better our chances (so long as we continue to think about what is best for the network stack as a whole and not what is best for network namespaces).

Eric

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers