Subject: net 2.6.25

Posted by den on Wed, 14 Nov 2007 09:32:10 GMT

View Forum Message <> Reply to Message

Hello, guys!

David has been opened net-2.6.25 tree yesterday, so we could send a new portion of NETNS there. If there are no objections, I'll send first six patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet & Unix sockets virtualization + some our /proc cleanups.

Regards, Den

Subject: Re: net 2.6.25

Posted by ebiederm on Wed, 14 Nov 2007 09:52:46 GMT

View Forum Message <> Reply to Message

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

> Hello, guys!

>

- > David has been opened net-2.6.25 tree yesterday, so we could send a new
- > portion of NETNS there. If there are no objections, I'll send first six
- > patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet &
- > Unix sockets virtualization + some our /proc cleanups.

I think I can pretty much agree on that set of patches. However I do strongly object to the form they are in at: http://git.openvz.org/?p=linux-2.6-netns

All of my subject lines (which are the primary summary of the patch) have been lost. So it is much more difficult to see what the patch is doing.

I will get some sleep and see what I can sort out.

Eric

Subject: Re: net 2.6.25

Posted by den on Wed, 14 Nov 2007 10:03:30 GMT

View Forum Message <> Reply to Message

Will you be happy, if I'll re-merge all of them from scratch preserving original subjects from the letters?

Regards,

## Den

```
Eric W. Biederman wrote:

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

> Hello, guys!

>> David has been opened net-2.6.25 tree yesterday, so we could send a new
>> portion of NETNS there. If there are no objections, I'll send first six
>> patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet &
>> Unix sockets virtualization + some our /proc cleanups.
>> I think I can pretty much agree on that set of patches. However I do strongly
> object to the form they are in at: http://git.openvz.org/?p=linux-2.6-netns
>> All of my subject lines (which are the primary summary of the patch) have
> been lost. So it is much more difficult to see what the patch is doing.
>> I will get some sleep and see what I can sort out.
>> Eric
> Eric
```

Subject: Re: net 2.6.25

Posted by Daniel Lezcano on Wed, 14 Nov 2007 10:31:27 GMT

View Forum Message <> Reply to Message

Denis V. Lunev wrote:

> Hello, guys!

>

- > David has been opened net-2.6.25 tree yesterday, so we could send a new
- > portion of NETNS there. If there are no objections, I'll send first six
- > patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet &
- > Unix sockets virtualization + some our /proc cleanups.

Did Pavel had time to look at RCU looking mechanism to replace rtnl\_lock/rtnl\_unlock to protect the netns list?

ps: isn't the openvz mailing list already subscribed to containers@

Subject: Re: net 2.6.25

Posted by den on Wed, 14 Nov 2007 10:52:55 GMT

View Forum Message <> Reply to Message

Daniel Lezcano wrote:

- > Denis V. Lunev wrote:
- >> Hello, guys!

>>

- >> David has been opened net-2.6.25 tree yesterday, so we could send a new
- >> portion of NETNS there. If there are no objections, I'll send first six
- >> patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet &
- >> Unix sockets virtualization + some our /proc cleanups.

>

- > Did Pavel had time to look at RCU looking mechanism to replace
- > rtnl\_lock/rtnl\_unlock to protect the netns list ? not yet
- > ps: isn't the openvz mailing list already subscribed to containers@ hmm:) Pavel say yes:) trying...

Containers mailing list

Containers@lists.linux-foundation.org

https://lists.linux-foundation.org/mailman/listinfo/containers

Subject: Re: Re: net 2.6.25

Posted by Pavel Emelianov on Wed, 14 Nov 2007 10:59:01 GMT

View Forum Message <> Reply to Message

Daniel Lezcano wrote:

- > Denis V. Lunev wrote:
- >> Hello, guys!

>>

- >> David has been opened net-2.6.25 tree yesterday, so we could send a new
- >> portion of NETNS there. If there are no objections, I'll send first six
- >> patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet &
- >> Unix sockets virtualization + some our /proc cleanups.

>

- > Did Pavel had time to look at RCU looking mechanism to replace
- > rtnl\_lock/rtnl\_unlock to protect the netns list ?

Unfortunately, I had not. Andrew has released the rc2-mm1 tree, so I'm now trying to prepare all the patches I wanted for this mm.

I will revisit this problem, but Denis said, that the problem we're trying to solve with the RCU protection can be solved w/o it, no?

> ps: isn't the openvz mailing list already subscribed to containers@

devel@openvz.org is.

Subject: Re: Re: net 2.6.25

Posted by Daniel Lezcano on Wed, 14 Nov 2007 11:05:09 GMT

View Forum Message <> Reply to Message

Pavel Emelyanov wrote:

- > Daniel Lezcano wrote:
- >> Denis V. Lunev wrote:
- >>> Hello, guys!

>>>

- >>> David has been opened net-2.6.25 tree yesterday, so we could send a new
- >>> portion of NETNS there. If there are no objections, I'll send first six
- >>> patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet &
- >>> Unix sockets virtualization + some our /proc cleanups.
- >> Did Pavel had time to look at RCU looking mechanism to replace
- >> rtnl\_lock/rtnl\_unlock to protect the netns list ?

>

- > Unfortunately, I had not. Andrew has released the rc2-mm1 tree, so
- > I'm now trying to prepare all the patches I wanted for this mm.

>

- > I will revisit this problem, but Denis said, that the problem we're
- > trying to solve with the RCU protection can be solved w/o it, no?

Yes, if we create a timer per network namespace for ipv6 garbage collecting, we can avoid to use a rcu, in this case.

I am doing that right now, but I am afraid we can fall into another case where we need to browse the network list and I thought having the RCU will bring more flexibility.

But, we can live without rcu and go ahead with ipv6 anyway:)

Containers mailing list

Containers@lists.linux-foundation.org

https://lists.linux-foundation.org/mailman/listinfo/containers

Subject: Re: net 2.6.25

Posted by ebiederm on Wed, 14 Nov 2007 16:03:30 GMT

View Forum Message <> Reply to Message

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

- > Daniel Lezcano wrote:
- >> Denis V. Lunev wrote:
- >>> Hello, guys!

>>>

- >>> David has been opened net-2.6.25 tree yesterday, so we could send a new
- >>> portion of NETNS there. If there are no objections, I'll send first six

>>> patches from OpenVZ GIT tomorrow, i.e. RTNL infrastructure, Packet &

>>> Unix sockets virtualization + some our /proc cleanups.

>>

>> Did Pavel had time to look at RCU looking mechanism to replace

>> rtnl\_lock/rtnl\_unlock to protect the netns list ?

> not yet

>

>> ps: isn't the openvz mailing list already subscribed to containers@

> hmm :) Pavel say yes :) trying...

Well honestly I think I would rather have a spin lock then an rcu thing.

It may not be a big deal because we already have to check the count and if it is zero ignore the namespace, so rcu may not add complications but changing something that already has interesting usage and lifetime rules to rcu sends up warning signs for me.

Eric

Containers mailing list
Containers@lists.linux-foundation.org

https://lists.linux-foundation.org/mailman/listinfo/containers

Subject: Re: net 2.6.25

Posted by ebiederm on Wed, 14 Nov 2007 18:19:38 GMT

View Forum Message <> Reply to Message

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

- > Will you be happy, if I'll re-merge all of them from scratch preserving
- > original subjects from the letters?

Yes that is my big issue.

Eric

Subject: Re: Re: net 2.6.25

Posted by kir on Thu, 15 Nov 2007 17:15:37 GMT

View Forum Message <> Reply to Message

Pavel Emelyanov wrote:

> Daniel Lezcano wrote:

```
>
>> Denis V. Lunev wrote:
>>
>
>> ps: isn't the openvz mailing list already subscribed to containers@
>>
>
> devel@openvz.org is.
>
```

That shouldn't stop you from posting to both -- mailman is smart enough to cope with two copies of the same message, so there will be no dupes.