Subject: current state of netns
Posted by den on Wed, 17 Oct 2007 11:27:06 GMT
View Forum Message <> Reply to Message

Hello, Eric!

I see that you quite busy and there is no reaction from Dave for your latest portion of netns patches. Right now, me and Pavel are working exclusively for mainstream.

May be we could bring a torch from your hands and start to push Dave Miller even with IPv4 staff. 3 weeks passed, no reaction for you latest code. Looks like it has been missed somehow... I even have to stop my fingers every day from touching a generic structures like flowi :)

Regards,
 Den

Subject: Re: current state of netns
Posted by ebiederm on Wed, 17 Oct 2007 18:16:53 GMT
View Forum Message <> Reply to Message

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

> Hello, Eric!
>
> I see that you quite busy and there is no reaction from Dave for your latest
> portion of netns patches. Right now, me and Pavel are working exclusively for
> mainstream.
>
> May be we could bring a torch from your hands and start to push Dave Miller even
> with IPv4 staff. 3 weeks passed, no reaction for you latest code. Looks like it
> has been missed somehow... I even have to stop my fingers every day from
> touching a generic structures like flowi :)

Short summary.
- The merge window opened late.
- All of the netns code needs to be to Dave Miller before the merge window.
- My last round of changes were not bug fixes and were sent after Dave
  had stopped accepting feature additions for 2.6.24

Therefore after the merge window when Dave Miller is ready to queue up more networking patches I expect progress can be made again.

I think the only thing that is happening is unfortunate timing.

I'm not really opposed to people taking my patches or something like

them cleaning them up and running with them, I just think the current
slow down bad timing. We have achieved the hard part which is to
get the core network namespace infrastructure accepted.

On another note. While I think using CONFIG_NET_NS is nice. I really
only introduced it so that production kernels can avoid enabling an
experimental feature. So far it still looks sane to me to remove
CONFIG_NET_NS when things are solid and we can remove the experimental
tag.

As for ipv4 and ipv6. However we do that we want to very carefully
sequence the patches so that we increasingly make the network
namespace infrastructure fine grained. Similar to make locks fine
grained. I did that for my core network namespaces patches but that
careful ordering still needs to happen for my ipv4 patches.

Eric

Subject: Re: current state of netns
Posted by Daniel Lezcano on Thu, 18 Oct 2007 08:57:35 GMT
View Forum Message <> Reply to Message

Eric W. Biederman wrote:
> "Denis V. Lunev" <den@sw.ru> writes:
>
>> Hello, Eric!
>>
>> I see that you quite busy and there is no reaction from Dave for your latest
>> portion of netns patches. Right now, me and Pavel are working exclusively for
>> mainstream.
>>
>> May be we could bring a torch from your hands and start to push Dave Miller even
>> with IPv4 staff. 3 weeks passed, no reaction for you latest code. Looks like it
>> has been missed somehow... I even have to stop my fingers every day from
>> touching a generic structures like flowi :)
>
> Short summary.
> - The merge window opened late.
> - All of the netns code needs to be to Dave Miller before the merge window.
> - My last round of changes were not bug fixes and were sent after Dave
>   had stopped accepting feature additions for 2.6.24
>
> Therefore after the merge window when Dave Miller is ready to queue up
> more networking patches I expect progress can be made again.
>
> I think the only thing that is happening is unfortunate timing.
>

> I'm not really opposed to people taking my patches or something like
> them cleaning them up and running with them, I just think the current
> slow down bad timing. We have achieved the hard part which is to
> get the core network namespace infrastructure accepted.
>
> On another note. While I think using CONFIG_NET_NS is nice. I really
> only introduced it so that production kernels can avoid enabling an
> experimental feature. So far it still looks sane to me to remove
> CONFIG_NET_NS when things are solid and we can remove the experimental
> tag.
>
> As for ipv4 and ipv6. However we do that we want to very carefully
> sequence the patches so that we increasingly make the network
> namespace infrastructure fine grained. Similar to make locks fine
> grained. I did that for my core network namespaces patches but that
> careful ordering still needs to happen for my ipv4 patches.

Denis, Pavel,

this is great to have you with us for netns. Do you mind if we follow
the rule : "patches sent to netdev@ are coming from Eric's git tree, any
enhancements are posted to Eric/containers" ?
So at least, we have the patches stacked and that give us time to review
and to test.

Eric, what do you think about that ?

By the way, Benjamin and I, we are making ipv6 per namespace. We will
send a first patchset for addrconf, ndisc, ip_fib6, fib6_rules probably
at the end of the week or at the begin of the next week.

We are also planning to choose a small patch subset from Eric's tree for
ipv4 to be proposed to containers@ before sending it to netdev@ (we
should be here very careful and send ipv4, piece by piece, and ensure at
all cost init_net_ns will not be broken).

I don't have a clear idea when the merge window will be closed. I guess,
we should resend af_netlink, af_unix and af_packet before sending
anything new, like af_inet.

Can we coordinate our effort, what do you plan to do ?

Regards.

 -- Daniel

Subject: Re: current state of netns
Posted by den on Thu, 18 Oct 2007 10:09:36 GMT
View Forum Message <> Reply to Message

Daniel Lezcano wrote:
> Eric W. Biederman wrote:
>> "Denis V. Lunev" <den@sw.ru> writes:
>>
>>> Hello, Eric!
>>>
>>> I see that you quite busy and there is no reaction from Dave for your
>>> latest
>>> portion of netns patches. Right now, me and Pavel are working
>>> exclusively for
>>> mainstream.
>>>
>>> May be we could bring a torch from your hands and start to push Dave
>>> Miller even
>>> with IPv4 staff. 3 weeks passed, no reaction for you latest code.
>>> Looks like it
>>> has been missed somehow... I even have to stop my fingers every day from
>>> touching a generic structures like flowi :)
>>
>> Short summary. - The merge window opened late.
>> - All of the netns code needs to be to Dave Miller before the merge
>> window.
>> - My last round of changes were not bug fixes and were sent after Dave
>>   had stopped accepting feature additions for 2.6.24
>>
>> Therefore after the merge window when Dave Miller is ready to queue up
>> more networking patches I expect progress can be made again.
>>
>> I think the only thing that is happening is unfortunate timing.
>>
>> I'm not really opposed to people taking my patches or something like
>> them cleaning them up and running with them, I just think the current
>> slow down bad timing. We have achieved the hard part which is to
>> get the core network namespace infrastructure accepted.
>>
>> On another note. While I think using CONFIG_NET_NS is nice. I really
>> only introduced it so that production kernels can avoid enabling an
>> experimental feature. So far it still looks sane to me to remove
>> CONFIG_NET_NS when things are solid and we can remove the experimental
>> tag.
>>
>> As for ipv4 and ipv6. However we do that we want to very carefully
>> sequence the patches so that we increasingly make the network
>> namespace infrastructure fine grained. Similar to make locks fine
>> grained. I did that for my core network namespaces patches but that

>> careful ordering still needs to happen for my ipv4 patches.
>
> Denis, Pavel,
>
> this is great to have you with us for netns. Do you mind if we follow
> the rule : "patches sent to netdev@ are coming from Eric's git tree, any
> enhancements are posted to Eric/containers" ?
> So at least, we have the patches stacked and that give us time to review
> and to test.
>
> Eric, what do you think about that ?

NETNS49 is outdated in comparison to mainstream now in respect to
locking and rtnl handling. Some other small features also counts.

I think a re-base should be scheduled for some time.

> By the way, Benjamin and I, we are making ipv6 per namespace. We will
> send a first patchset for addrconf, ndisc, ip_fib6, fib6_rules probably
> at the end of the week or at the begin of the next week.
>
> We are also planning to choose a small patch subset from Eric's tree for
> ipv4 to be proposed to containers@ before sending it to netdev@ (we
> should be here very careful and send ipv4, piece by piece, and ensure at
> all cost init_net_ns will not be broken).
>
> I don't have a clear idea when the merge window will be closed. I guess,
> we should resend af_netlink, af_unix and af_packet before sending
> anything new, like af_inet.

IMHO we should update them to new rtnl and resend.

> Can we coordinate our effort, what do you plan to do ?

Right now there is no concrete plan. We thought about IPv6 also. But, as
I see two major areas
- prepare finegrained pieces for mainstream
- re-base current infrastructure

Making IPv6 in NETNS49 does not seems good to me. Serious amount of job
will be lost... Kernel is changed even by us :) And even _NAMESPACE_
code is changed during committing process.

At least some areas for generic optimizations are revealed.

Regards,
 Den

Subject: Re: current state of netns
Posted by Daniel Lezcano on Thu, 18 Oct 2007 10:37:57 GMT

Denis V. Lunev wrote:
> Daniel Lezcano wrote:
>> Eric W. Biederman wrote:
>>> "Denis V. Lunev" <den@sw.ru> writes:
>>>
>>>> Hello, Eric!
>>>>
>>>> I see that you quite busy and there is no reaction from Dave for
>>>> your latest
>>>> portion of netns patches. Right now, me and Pavel are working
>>>> exclusively for
>>>> mainstream.
>>>>
>>>> May be we could bring a torch from your hands and start to push Dave
>>>> Miller even
>>>> with IPv4 staff. 3 weeks passed, no reaction for you latest code.
>>>> Looks like it
>>>> has been missed somehow... I even have to stop my fingers every day
>>>> from
>>>> touching a generic structures like flowi :)
>>>
>>> Short summary. - The merge window opened late.
>>> - All of the netns code needs to be to Dave Miller before the merge
>>> window.
>>> - My last round of changes were not bug fixes and were sent after Dave
>>>   had stopped accepting feature additions for 2.6.24
>>>
>>> Therefore after the merge window when Dave Miller is ready to queue up
>>> more networking patches I expect progress can be made again.
>>>
>>> I think the only thing that is happening is unfortunate timing.
>>>
>>> I'm not really opposed to people taking my patches or something like
>>> them cleaning them up and running with them, I just think the current
>>> slow down bad timing. We have achieved the hard part which is to
>>> get the core network namespace infrastructure accepted.
>>>
>>> On another note. While I think using CONFIG_NET_NS is nice. I really
>>> only introduced it so that production kernels can avoid enabling an
>>> experimental feature. So far it still looks sane to me to remove
>>> CONFIG_NET_NS when things are solid and we can remove the experimental
>>> tag.
>>>
>>> As for ipv4 and ipv6. However we do that we want to very carefully
>>> sequence the patches so that we increasingly make the network

>>> namespace infrastructure fine grained. Similar to make locks fine
>>> grained. I did that for my core network namespaces patches but that
>>> careful ordering still needs to happen for my ipv4 patches.
>>
>> Denis, Pavel,
>>
>> this is great to have you with us for netns. Do you mind if we follow
>> the rule : "patches sent to netdev@ are coming from Eric's git tree,
>> any enhancements are posted to Eric/containers" ?
>> So at least, we have the patches stacked and that give us time to
>> review and to test.
>>
>> Eric, what do you think about that ?
>
> NETNS49 is outdated in comparison to mainstream now in respect to
> locking and rtnl handling. Some other small features also counts.
>
> I think a re-base should be scheduled for some time.

Well, I think Eric is quite busy. But he will certainly rebase netns49
very soon. I feel much more confortable to have Eric federating the patches.

>
>> By the way, Benjamin and I, we are making ipv6 per namespace. We will
>> send a first patchset for addrconf, ndisc, ip_fib6, fib6_rules
>> probably at the end of the week or at the begin of the next week.
>>
>> We are also planning to choose a small patch subset from Eric's tree
>> for ipv4 to be proposed to containers@ before sending it to netdev@
>> (we should be here very careful and send ipv4, piece by piece, and
>> ensure at all cost init_net_ns will not be broken).
>>
>> I don't have a clear idea when the merge window will be closed. I
>> guess, we should resend af_netlink, af_unix and af_packet before
>> sending anything new, like af_inet.
>
> IMHO we should update them to new rtnl and resend.
>
>> Can we coordinate our effort, what do you plan to do ?
>
> Right now there is no concrete plan. We thought about IPv6 also. But, as
> I see two major areas
> - prepare finegrained pieces for mainstream
> - re-base current infrastructure
> Making IPv6 in NETNS49 does not seems good to me. Serious amount of job
> will be lost... Kernel is changed even by us :) And even _NAMESPACE_
> code is changed during committing process.

That's the life :)

By the way, you changes are not functional changes, you patches are code factorization, fixes and optimization, no ? So that will not change the global approach for the ipv6 namespace.

Rebasing netns on the latest Dave Miller tree is obviously mandatory. But what do you expect to do with the rebase ? Do you will create a git repository at openvz.org ? I just don't want to have a fork of netns, that will be a total mess and invalidate all testing done on it.

Eric was doing the rebasing since the beginning and federating the patches. When something is posted to netdev@, it is a little piece of netns, so rebasing this to the latest Dave Miller tree is quite easy.

But, in the other hand, you are making a very good job with netns and I don't want to be push back. So, if I understand correctly, you wish to have a netns released more often in order to reduce the delta between Dave tree and netns tree, right ?

> At least some areas for generic optimizations are revealed.