Subject: Re: Network virtualization/isolation
Posted by ebiederm on Sun, 03 Dec 2006 16:00:59 GMT
View Forum Message <> Reply to Message

Ok. Just a quick summary of where I see the discussion.

We all agree that L2 isolation is needed at some point.

The approaches discussed for L2 and L3 are sufficiently orthogonal
that we can implement then in either order. You would need to
unshare L3 to unshare L2, but if we think of them as two separate
namespaces we are likely to be in better shape.

The L3 discussion still has the problem that there has not been
agreement on all of the semantics yet.

More comments after I get some sleep.

Eric

Subject: Re: Network virtualization/isolation
Posted by Mishin Dmitry on Mon, 04 Dec 2006 15:19:00 GMT
View Forum Message <> Reply to Message

On Sunday 03 December 2006 19:00, Eric W. Biederman wrote:
> Ok. Just a quick summary of where I see the discussion.
>
> We all agree that L2 isolation is needed at some point.
As we all agreed on this, may be it is time to send patches one-by-one?
For the beggining, I propose to resend Cedric's empty namespace patch as base
for others - it is really empty, but necessary in order to move further.

After this patch and the following net namespace unshare patch will be
accepted, I could send network devices virtualization patches for review and
discussion.

What do you think?

>
> The approaches discussed for L2 and L3 are sufficiently orthogonal
> that we can implement then in either order. You would need to
> unshare L3 to unshare L2, but if we think of them as two separate
> namespaces we are likely to be in better shape.
>
> The L3 discussion still has the problem that there has not been
> agreement on all of the semantics yet.
>

> More comments after I get some sleep.
>
> Eric
> -
> To unsubscribe from this list: send the line "unsubscribe netdev" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html

--
Thanks,
Dmitry.

Subject: Re: Network virtualization/isolation
Posted by ebiederm on Mon, 04 Dec 2006 15:45:42 GMT
View Forum Message <> Reply to Message

Dmitry Mishin <dim@openvz.org> writes:

> On Sunday 03 December 2006 19:00, Eric W. Biederman wrote:
>> Ok. Just a quick summary of where I see the discussion.
>>
>> We all agree that L2 isolation is needed at some point.
> As we all agreed on this, may be it is time to send patches one-by-one?
> For the beggining, I propose to resend Cedric's empty namespace patch as base
> for others - it is really empty, but necessary in order to move further.
>
> After this patch and the following net namespace unshare patch will be
> accepted, I could send network devices virtualization patches for review and
> discussion.
>
> What do you think?

I think sending out these patches for review sounds great.

For merge order I think enabling the unshare/clone flags to anyone
but developers should be about the last thing we do.

Starting with clone/unshare sounds to me like hitching up the cart
before it is built.

I really need to focus on finishing up the pid namespace, so except
for a little review and conversation I'm not going to help much on the
network side.

Of course I need to mess with unix domain sockets to properly
implement the pid namespace. Because of the pid credential passing.

Eric

Subject: Re: Network virtualization/isolation
Posted by Herbert Poetzl on Mon, 04 Dec 2006 16:43:32 GMT
View Forum Message <> Reply to Message

On Mon, Dec 04, 2006 at 06:19:00PM +0300, Dmitry Mishin wrote:
> On Sunday 03 December 2006 19:00, Eric W. Biederman wrote:
> > Ok. Just a quick summary of where I see the discussion.
> >
> > We all agree that L2 isolation is needed at some point.

> As we all agreed on this, may be it is time to send patches
> one-by-one? For the beggining, I propose to resend Cedric's
> empty namespace patch as base for others - it is really empty,
> but necessary in order to move further.
>
> After this patch and the following net namespace unshare
> patch will be accepted,

well, I have neither seen any performance tests showing
that the following is true:

- no change on network performance without the
  space enabled
- no change on network performance on the host
  with the network namespaces enabled
- no measureable overhead inside the network
  namespace
- good scaleability for a larger number of network
  namespaces

> I could send network devices virtualization patches for
> review and discussion.

that won't hurt ...

best,
Herbert

> What do you think?
>
> > The approaches discussed for L2 and L3 are sufficiently orthogonal
> > that we can implement then in either order. You would need to
> > unshare L3 to unshare L2, but if we think of them as two separate
> > namespaces we are likely to be in better shape.
> >

> > The L3 discussion still has the problem that there has not been
> > agreement on all of the semantics yet.
> >
> > More comments after I get some sleep.
> >
> > Eric
> > -
> > To unsubscribe from this list: send the line "unsubscribe netdev" in
> > the body of a message to majordomo@vger.kernel.org
> > More majordomo info at http://vger.kernel.org/majordomo-info.html
>
> --
> Thanks,
> Dmitry.
> _____
> Containers mailing list
> Containers@lists.osdl.org
> https://lists.osdl.org/mailman/listinfo/containers

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: Network virtualization/isolation
Posted by Mishin Dmitry on Mon, 04 Dec 2006 17:02:48 GMT
View Forum Message <> Reply to Message

On Monday 04 December 2006 19:43, Herbert Poetzl wrote:
> On Mon, Dec 04, 2006 at 06:19:00PM +0300, Dmitry Mishin wrote:
> > On Sunday 03 December 2006 19:00, Eric W. Biederman wrote:
> > > Ok. Just a quick summary of where I see the discussion.
> > >
> > > We all agree that L2 isolation is needed at some point.
> >
> > As we all agreed on this, may be it is time to send patches
> > one-by-one? For the beggining, I propose to resend Cedric's
> > empty namespace patch as base for others - it is really empty,
> > but necessary in order to move further.
> >
> > After this patch and the following net namespace unshare
> > patch will be accepted,
>
> well, I have neither seen any performance tests showing
> that the following is true:
>
> - no change on network performance without the
>   space enabled

> - no change on network performance on the host
> with the network namespaces enabled
> - no measureable overhead inside the network
> namespace
> - good scaleability for a larger number of network
> namespaces
These questions are for complete L2 implementation, not for these 2 empty patches. If you need some data relating to Andrey's implementation, I'll get it. Which test do you accept?

>
> > I could send network devices virtualization patches for
> > review and discussion.
>
> that won't hurt ...
>
> best,
> Herbert
>
> > What do you think?
> >
> > > The approaches discussed for L2 and L3 are sufficiently orthogonal
> > > that we can implement then in either order. You would need to
> > > unshare L3 to unshare L2, but if we think of them as two separate
> > > namespaces we are likely to be in better shape.
> > >
> > > The L3 discussion still has the problem that there has not been
> > > agreement on all of the semantics yet.
> > >
> > > More comments after I get some sleep.
> > >
> > > Eric
> > > -
> > > To unsubscribe from this list: send the line "unsubscribe netdev" in
> > > the body of a message to majordomo@vger.kernel.org
> > > More majordomo info at http://vger.kernel.org/majordomo-info.html
> >
> > --
> > Thanks,
> > Dmitry.
> > _____
> > Containers mailing list
> > Containers@lists.osdl.org
> > https://lists.osdl.org/mailman/listinfo/containers

--
Thanks,
Dmitry.

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers