Subject: Re: [patch 0/1] [RFC][net namespace] veth ioctl management
Posted by ebiederm on Mon, 19 Feb 2007 20:01:21 GMT
View Forum Message <> Reply to Message

Dmitry Mishin <dim@openvz.org> writes:

> Fully agree. But as I can see, your code arises no more comments, than ours.
> So, we need to find other ways. Do you have more ideas?

Yes.

To some extent we should probably compare notes and see which parts
of the various implementations are good/bad.

For the most part from what I could see at least when doing L2 level
work the two patchsets touched roughly the same code in roughly the
same ways the differences the big differences being in how complete
one patchset one area or not. So while push/pop helped a little with
the argument passing it was small enough it wasn't a big deal either
way.

The planetlab folks are busily evaluating and collecting some benchmarks
numbers. Last I heard OpenVZ, vs mine, vs native were all pretty much
a wash on bandwidth and latency. For cpu consumption OpenVZ and mine
when multiple guests were running were worse by a small factor, cache
effects was the guess. If those results hold and the costs of an L2
namespace stays firmly in the noise it will be hard to justify any
kind of L3 namespace.

My sysctl stuff has gone in, and I will have sysfs support as soon
as the network sysfs support settles down. So there is some progress
there.

I suspect we won't have any real problems merging an tunnel device
like etun or veth as long as we don't need the push/pop in the middle
to make it work. I was thinking about that a little.

I asked David Miller if he had looked at what I had posted and he
replied that he had wanted to but he was swamped with bug fixes
and sparc maintenance.

So I expect what happened is that is the I posted too much code at
once so it was hard to digest.

My current plan is:
- kill the stupid irq migration bug on x86_64 (sucks way to much time)
- finish up the sysctl and other network namespace helper support

- discuss my network namespace patches and see what Dmitry and
  Daniel and any other interested parties think of them.
- merge a tunnel device
- post network namespace code in the smallest chunk I can stand
  and ask that it be included.

  Hopefully real working real working code that is ready to go
  will either get merged or there will be reasonable feedback
  on why it was not merged.

- Somewhere in there general maintenance, testing and completing
  of the network namespace code I have.

Eric

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

Subject: Re: [patch 0/1] [RFC][net namespace] veth ioctl management
Posted by Daniel Lezcano on Tue, 20 Feb 2007 09:43:55 GMT
View Forum Message <> Reply to Message

Eric W. Biederman wrote:
> Dmitry Mishin <dim@openvz.org> writes:
>
>
>> Fully agree. But as I can see, your code arises no more comments, than ours.
>> So, we need to find other ways. Do you have more ideas?
>
> Yes.
>
> To some extent we should probably compare notes and see which parts
> of the various implementations are good/bad.
>
> For the most part from what I could see at least when doing L2 level
> work the two patchsets touched roughly the same code in roughly the
> same ways the differences the big differences being in how complete
> one patchset one area or not. So while push/pop helped a little with
> the argument passing it was small enough it wasn't a big deal either
> way.
>
> The planetlab folks are busily evaluating and collecting some benchmarks
> numbers. Last I heard OpenVZ, vs mine, vs native were all pretty much
> a wash on bandwidth and latency. For cpu consumption OpenVZ and mine
> when multiple guests were running were worse by a small factor, cache
> effects was the guess. If those results hold and the costs of an L2

> namespace stays firmly in the noise it will be hard to justify any
> kind of L3 namespace.

Definitively, you are against L3 namespace :)

> My sysctl stuff has gone in, and I will have sysfs support as soon
> as the network sysfs support settles down. So there is some progress
> there.
>
> I suspect we won't have any real problems merging an tunnel device
> like etun or veth as long as we don't need the push/pop in the middle
> to make it work. I was thinking about that a little.
>
> I asked David Miller if he had looked at what I had posted and he
> replied that he had wanted to but he was swamped with bug fixes
> and sparc maintenance.
>
> So I expect what happened is that is the I posted too much code at
> once so it was hard to digest.

Yes.

> My current plan is:
> - kill the stupid irq migration bug on x86_64 (sucks way to much time)
> - finish up the sysctl and other network namespace helper support
> - discuss my network namespace patches and see what Dmitry and
>   Daniel and any other interested parties think of them.

Ok for me, please send part by part patches, it will be more digest.

> - merge a tunnel device
> - post network namespace code in the smallest chunk I can stand
>   and ask that it be included.

>   Hopefully real working real working code that is ready to go
>   will either get merged or there will be reasonable feedback
>   on why it was not merged.

I agree.

> - Somewhere in there general maintenance, testing and completing
>   of the network namespace code I have.

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

Subject: Re: [patch 0/1] [RFC][net namespace] veth ioctl management
Posted by ebiederm on Tue, 20 Feb 2007 17:52:12 GMT
View Forum Message <> Reply to Message

Daniel Lezcano <dlezcano@fr.ibm.com> writes:

>> My current plan is:
>> - kill the stupid irq migration bug on x86_64 (sucks way to much time)
>> - finish up the sysctl and other network namespace helper support
>> - discuss my network namespace patches and see what Dmitry and
>> Daniel and any other interested parties think of them.
>
> Ok for me, please send part by part patches, it will be more digest.

I can only send my patches slower. My patches were well factored.
The volume and size of the patches reflects the scope of the
problem.

One of the patches I sent was my etun driver. Please take a second
look at what I posted.

Not that you really need an example of a networking driver that uses
sysfs to see how that can be done.

Eric

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