
Subject: containers and cgroups mini-summit @ Linux Plumbers

Posted by [kir](#) on Wed, 11 Jul 2012 21:41:53 GMT

[View Forum Message](#) <> [Reply to Message](#)

Gentlemen,

We are organizing containers mini-summit during next Linux Plumbers (San Diego, August 29-31).

The idea is to gather and discuss everything relevant to namespaces, cgroups, resource management, checkpoint-restore and so on.

We are trying to come up with a list of topics to discuss, so please reply with topic suggestions, and let me know if you are going to come.

I probably forgot a few more people (such as, I am not sure who else from Google is working on cgroups stuff), so feel free to forward this to anyone you believe should go, or just let me know whom I missed.

Regards,
Kir.

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [Serge E. Hallyn](#) on Thu, 12 Jul 2012 03:47:41 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Kir Kolyshkin (kir@openvz.org):

> Gentlemen,

>

> We are organizing containers mini-summit during next Linux Plumbers
> (San Diego, August 29-31).

> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.

> We are trying to come up with a list of topics to discuss, so please
> reply with topic suggestions, and
> let me know if you are going to come.

I won't be there. I hope it'll be possible to call in remotely one way or another.

Topics I'm most interested in right now, beside completion of user namespace (if anything needs to be discussed there), is syslog

namespace and what to do about loop devices and udev/uevents (devices namespace?). Checkpoint/restart is moving along nicely, perhaps only a progress update and discussion of any limitations?

> I probably forgot a few more people (such as, I am not sure who else
> from Google is working
> on cgroups stuff), so feel free to forward this to anyone you
> believe should go,
> or just let me know whom I missed.
>
> Regards,
> Kir.

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Frederic Weisbecker](#) on Thu, 12 Jul 2012 13:26:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Jul 12, 2012 at 01:41:53AM +0400, Kir Kolyshkin wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers
> (San Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.
>
> We are trying to come up with a list of topics to discuss, so please
> reply with topic suggestions, and
> let me know if you are going to come.

I wish I could, but this year's LPC date conflict with my holidays.

Sorry.

Thanks.

>
> I probably forgot a few more people (such as, I am not sure who else
> from Google is working
> on cgroups stuff), so feel free to forward this to anyone you
> believe should go,
> or just let me know whom I missed.
>
> Regards,
> Kir.

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [kir](#) on Fri, 13 Jul 2012 15:09:07 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 07/12/2012 07:47 AM, Serge Hallyn wrote:

> Quoting Kir Kolyshkin (kir@openvz.org):

>> Gentlemen,

>>

>> We are organizing containers mini-summit during next Linux Plumbers

>> (San Diego, August 29-31).

>> The idea is to gather and discuss everything relevant to namespaces,

>> cgroups, resource management,

>> checkpoint-restore and so on.

>> We are trying to come up with a list of topics to discuss, so please

>> reply with topic suggestions, and

>> let me know if you are going to come.

> I won't be there. I hope it'll be possible to call in remotely one way or

> another.

We will try to set up something (Skype? Google+ Hangout? A conventional phone line

from event organizers?) but from some previous experience it can't be very successful

if there are more than about 10 people in the room.

> Topics I'm most interested in right now, beside completion of user

> namespace (if anything needs to be discussed there), is syslog

> namespace and what to do about loop devices and udev/uevents (devices

> namespace?). Checkpoint/restart is moving along nicely, perhaps only

> a progress update and discussion of any limitations?

Thanks, will add it to the list.

>

>> I probably forgot a few more people (such as, I am not sure who else

>> from Google is working

>> on cgroups stuff), so feel free to forward this to anyone you

>> believe should go,

>> or just let me know whom I missed.

>>

>> Regards,

>> Kir.

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [Serge E. Hallyn](#) on Mon, 16 Jul 2012 20:29:50 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Kir Kolyshkin (kir@openvz.org):

> On 07/12/2012 07:47 AM, Serge Hallyn wrote:

> >Quoting Kir Kolyshkin (kir@openvz.org):

> >>Gentlemen,

> >>

> >>We are organizing containers mini-summit during next Linux Plumbers

> >>(San Diego, August 29-31).

> >>The idea is to gather and discuss everything relevant to namespaces,

> >>cgroups, resource management,

> >>checkpoint-restore and so on.

> >>We are trying to come up with a list of topics to discuss, so please

> >>reply with topic suggestions, and

> >>let me know if you are going to come.

> >I won't be there. I hope it'll be possible to call in remotely one way or

> >another.

>

> We will try to set up something (Skype? Google+ Hangout? A

> conventional phone line

> from event organizers?) but from some previous experience it can't

> be very successful

> if there are more than about 10 people in the room.

Yeah, I remember we tried it in 2007 or so...

hangouts crash every computer I own. I'm fine with skype or a phone line.

> >Topics I'm most interested in right now, beside completion of user

> >namespace (if anything needs to be discussed there), is syslog

> >namespace and what to do about loop devices and udev/uevents (devices

> >namespace?). Checkpoint/restart is moving along nicely, perhaps only

> >a progress update and discussion of any limitations?

>

> Thanks, will add it to the list.

Great, thanks.

> >>I probably forgot a few more people (such as, I am not sure who else

> >>from Google is working

> >>on cgroups stuff), so feel free to forward this to anyone you

> >>believe should go,

> >>or just let me know whom I missed.

> >>

> >>Regards,

> >> Kir.

>

>

-serge

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Dhaval Giani](#) on Mon, 16 Jul 2012 21:08:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

On Wed, Jul 11, 2012 at 5:41 PM, Kir Kolyshkin <kir@openvz.org> wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers (San
> Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.
>
> We are trying to come up with a list of topics to discuss, so please reply
> with topic suggestions, and
> let me know if you are going to come.
>

I will be at LPC, so I should be able to attend.

> I probably forgot a few more people (such as, I am not sure who else from
> Google is working
> on cgroups stuff), so feel free to forward this to anyone you believe should
> go,
> or just let me know whom I missed.

I would be interested mainly in discussing how cgroups core is changing, imposing a set of rules on how new subsystems are being designed (hierarchical vs non-hierarchical, and all that mess).

Dhaval

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Balbir Singh](#) on Tue, 17 Jul 2012 06:59:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Jul 12, 2012 at 3:11 AM, Kir Kolyshkin <kir@openvz.org> wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers (San
> Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.
>

> We are trying to come up with a list of topics to discuss, so please reply
> with topic suggestions, and
> let me know if you are going to come.
>
> I probably forgot a few more people (such as, I am not sure who else from
> Google is working
> on cgroups stuff), so feel free to forward this to anyone you believe should
> go,
> or just let me know whom I missed.

I can be and would love to be there if I get an invite and get to know in advance. I'll need to work out my travel plans/budget to see if I can get an approval. I think there are other interesting cases for containers especially on new generation, new form factor devices.

Thanks,
Balbir

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Daniel Lezcano](#) on Tue, 17 Jul 2012 07:26:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 07/11/2012 11:41 PM, Kir Kolyshkin wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers (San
> Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.
>
> We are trying to come up with a list of topics to discuss, so please
> reply with topic suggestions, and
> let me know if you are going to come.
>
> I probably forgot a few more people (such as, I am not sure who else
> from Google is working
> on cgroups stuff), so feel free to forward this to anyone you believe
> should go,
> or just let me know whom I missed.

Hi Kir,

I have a presentation for LPC and I am awaiting the approval for the funding. If it is accepted I will be there.

One point to address could be the time virtualization.

Thanks
-- Daniel

--

Follow Linaro: <<http://www.facebook.com/pages/Linaro>> Facebook |
<<http://twitter.com/#!/linaroorg>> Twitter |
<<http://www.linaro.org/linaro-blog/>> Blog

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Tejun Heo](#) on Tue, 17 Jul 2012 20:06:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

I think I can be there.

Thanks.

--

tejun

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Paul Turner](#) on Tue, 17 Jul 2012 20:15:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Jul 17, 2012 at 1:06 PM, Tejun Heo <tj@kernel.org> wrote:

> Hello,

>

> I think I can be there.

>

> Thanks.

>

> --

> tejun

Likewise,

- Paul

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [James Bottomley](#) on Wed, 18 Jul 2012 10:36:08 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, 2012-07-17 at 12:29 +0530, Balbir Singh wrote:

> On Thu, Jul 12, 2012 at 3:11 AM, Kir Kolyshkin <kir@openvz.org> wrote:

> > Gentlemen,

> >

> > We are organizing containers mini-summit during next Linux Plumbers (San

> > Diego, August 29-31).

> > The idea is to gather and discuss everything relevant to namespaces,

> > cgroups, resource management,

> > checkpoint-restore and so on.

> >

> > We are trying to come up with a list of topics to discuss, so please reply

> > with topic suggestions, and

> > let me know if you are going to come.

> >

> > I probably forgot a few more people (such as, I am not sure who else from

> > Google is working

> > on cgroups stuff), so feel free to forward this to anyone you believe should

> > go,

> > or just let me know whom I missed.

>

> I can be and would love to be there if I get an invite and get to know

> in advance. I'll need to work out my travel plans/budget to see if I

> can get an approval. I think there are other interesting cases for

> containers especially on new generation, new form factor devices.

You don't need an invite: This is Plumbers, not the kernel summit.

Plumbers Mini Summits/Micro Conferences are open to all registered

Plumbers attendees.

James

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [Johannes Weiner](#) on Fri, 20 Jul 2012 16:30:29 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Thu, Jul 12, 2012 at 01:41:53AM +0400, Kir Kolyshkin wrote:

> Gentlemen,

>

> We are organizing containers mini-summit during next Linux Plumbers

> (San Diego, August 29-31).

> The idea is to gather and discuss everything relevant to namespaces,

> cgroups, resource management,

> checkpoint-restore and so on.

>

> We are trying to come up with a list of topics to discuss, so please
> reply with topic suggestions, and
> let me know if you are going to come.

I'll be there.

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [KAMEZAWA Hiroyuki](#) on Sat, 21 Jul 2012 01:59:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

(2012/07/12 6:41), Kir Kolyshkin wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers (San Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces, cgroups, resource management,
> checkpoint-restore and so on.
>
> We are trying to come up with a list of topics to discuss, so please reply with topic suggestions, and
> let me know if you are going to come.
>

I think I'll be there.

-Kame

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Wed, 25 Jul 2012 08:48:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 07/17/2012 11:26 AM, Daniel Lezcano wrote:

> On 07/11/2012 11:41 PM, Kir Kolyshkin wrote:
>> Gentlemen,
>>
>> We are organizing containers mini-summit during next Linux Plumbers (San
>> Diego, August 29-31).
>> The idea is to gather and discuss everything relevant to namespaces,
>> cgroups, resource management,
>> checkpoint-restore and so on.
>>
>> We are trying to come up with a list of topics to discuss, so please
>> reply with topic suggestions, and
>> let me know if you are going to come.

>>
>> I probably forgot a few more people (such as, I am not sure who else
>> from Google is working
>> on cgroups stuff), so feel free to forward this to anyone you believe
>> should go,
>> or just let me know whom I missed.

>
> Hi Kir,
>
> I have a presentation for LPC and I am awaiting the approval for the
> funding. If it is accepted I will be there.
>
> One point to address could be the time virtualization.
>

What exactly do you have in mind for that ?

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Wed, 25 Jul 2012 08:55:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers (San
> Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.
>
> We are trying to come up with a list of topics to discuss, so please
> reply with topic suggestions, and
> let me know if you are going to come.
>
> I probably forgot a few more people (such as, I am not sure who else
> from Google is working
> on cgroups stuff), so feel free to forward this to anyone you believe
> should go,
> or just let me know whom I missed.
>
> Regards,
> Kir.

BTW, sorry for not replying before (vacations + post-vacations laziness)

I would be interested in adding /proc virtualization to the discussion.
By now it seems userspace would be the best place for that to happen, in

a fuse overlay. I know Daniel has an initial implementation of that, and it would be good to have it as library that both OpenVZ and LXC (and whoever else wants) can use.

Shouldn't take much time...

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [ebiederm](#) on Wed, 25 Jul 2012 10:00:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Glauber Costa <glommer@parallels.com> writes:

> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:
>> Gentlemen,
>>
>> We are organizing containers mini-summit during next Linux Plumbers (San
>> Diego, August 29-31).
>> The idea is to gather and discuss everything relevant to namespaces,
>> cgroups, resource management,
>> checkpoint-restore and so on.
>>
>> We are trying to come up with a list of topics to discuss, so please
>> reply with topic suggestions, and
>> let me know if you are going to come.
>>
>> I probably forgot a few more people (such as, I am not sure who else
>> from Google is working
>> on cgroups stuff), so feel free to forward this to anyone you believe
>> should go,
>> or just let me know whom I missed.
>>
>> Regards,
>> Kir.
>
> BTW, sorry for not replying before (vacations + post-vacations laziness)
>
> I would be interested in adding /proc virtualization to the discussion.
> By now it seems userspace would be the best place for that to happen, in
> a fuse overlay. I know Daniel has an initial implementation of that, and
> it would be good to have it as library that both OpenVZ and LXC (and
> whoever else wants) can use.
>
> Shouldn't take much time...

What would you need proc virtualization for?

Eric

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Wed, 25 Jul 2012 10:00:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 07/25/2012 02:00 PM, Eric W. Biederman wrote:

> Glauber Costa <glommer@parallels.com> writes:

>

>> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:

>>> Gentlemen,

>>>

>>> We are organizing containers mini-summit during next Linux Plumbers (San

>>> Diego, August 29-31).

>>> The idea is to gather and discuss everything relevant to namespaces,

>>> cgroups, resource management,

>>> checkpoint-restore and so on.

>>>

>>> We are trying to come up with a list of topics to discuss, so please

>>> reply with topic suggestions, and

>>> let me know if you are going to come.

>>>

>>> I probably forgot a few more people (such as, I am not sure who else

>>> from Google is working

>>> on cgroups stuff), so feel free to forward this to anyone you believe

>>> should go,

>>> or just let me know whom I missed.

>>>

>>> Regards,

>>> Kir.

>>

>> BTW, sorry for not replying before (vacations + post-vacations laziness)

>>

>> I would be interested in adding /proc virtualization to the discussion.

>> By now it seems userspace would be the best place for that to happen, in

>> a fuse overlay. I know Daniel has an initial implementation of that, and

>> it would be good to have it as library that both OpenVZ and LXC (and

>> whoever else wants) can use.

>>

>> Shouldn't take much time...

>

> What would you need proc virtualization for?

>

proc provides a lot of information that userspace tools rely upon.

For instance, when running top, you can draw per-process figures from what we have now, but you can't make sense of percentages without aggregating container-wide information.

When you read /proc/cpuinfo, as well, you would expect to see something that matches your container configuration.

"free" is another example. The list go on.

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Wed, 25 Jul 2012 10:02:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 07/25/2012 02:00 PM, Glauber Costa wrote:

> On 07/25/2012 02:00 PM, Eric W. Biederman wrote:

>> Glauber Costa <glommer@parallels.com> writes:

>>

>>> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:

>>>> Gentlemen,

>>>>

>>>> We are organizing containers mini-summit during next Linux Plumbers (San
>>>> Diego, August 29-31).

>>>> The idea is to gather and discuss everything relevant to namespaces,

>>>> cgroups, resource management,

>>>> checkpoint-restore and so on.

>>>>

>>>> We are trying to come up with a list of topics to discuss, so please

>>>> reply with topic suggestions, and

>>>> let me know if you are going to come.

>>>>

>>>> I probably forgot a few more people (such as, I am not sure who else

>>>> from Google is working

>>>> on cgroups stuff), so feel free to forward this to anyone you believe

>>>> should go,

>>>> or just let me know whom I missed.

>>>>

>>>> Regards,

>>>> Kir.

>>>

>>> BTW, sorry for not replying before (vacations + post-vacations laziness)

>>>

>>> I would be interested in adding /proc virtualization to the discussion.

>>> By now it seems userspace would be the best place for that to happen, in

>>> a fuse overlay. I know Daniel has an initial implementation of that, and

>>> it would be good to have it as library that both OpenVZ and LXC (and

>>> whoever else wants) can use.

>>>

>>> Shouldn't take much time...

>>

>> What would you need proc virtualization for?

>>

>

> proc provides a lot of information that userspace tools rely upon.

> For instance, when running top, you can draw per-process figures from
> what we have now, but you can't make sense of percentages without
> aggregating container-wide information.
>
> When you read /proc/cpuinfo, as well, you would expect to see something
> that matches your container configuration.
>
> "free" is another example. The list go on.
>
>

Also, Eric: Will you be around LPC? Would you be willing to hold a session about pid/user namespaces ?

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Wed, 25 Jul 2012 10:53:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers (San
> Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.
>
> We are trying to come up with a list of topics to discuss, so please
> reply with topic suggestions, and
> let me know if you are going to come.
>
> I probably forgot a few more people (such as, I am not sure who else
> from Google is working
> on cgroups stuff), so fill free to forward this to anyone you believe
> should go,
> or just let me know whom I missed.
>
> Regards,
> Kir.

I just came up with the following preliminary list of sessions:

<http://wiki.linuxplumbersconf.org/2012:containers>

Since people mostly said what they wanted to talk about, but without extensive descriptions, I took the liberty of coming up with a small text for each in the blueprints. If you believe this is inaccurate, or

would like to see it extended (although I personally don't see the point about going into very formal and deep details here), just let me know and I will edit it.

This is all still subject to change.

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [ebiederm](#) on Wed, 25 Jul 2012 11:16:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

Glauber Costa <glommer@parallels.com> writes:

> Also, Eric: Will you be around LPC? Would you be willing to hold a
> session about pid/user namespaces ?

Tentatively. My plans are not fixed in stone but I have some intention of figuring that out.

Eric

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Serge E. Hallyn](#) on Wed, 25 Jul 2012 11:51:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Glauber Costa (glommer@parallels.com):

> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:

>> Gentlemen,

>>

>> We are organizing containers mini-summit during next Linux Plumbers (San
>> Diego, August 29-31).

>> The idea is to gather and discuss everything relevant to namespaces,
>> cgroups, resource management,
>> checkpoint-restore and so on.

>>

>> We are trying to come up with a list of topics to discuss, so please
>> reply with topic suggestions, and
>> let me know if you are going to come.

>>

>> I probably forgot a few more people (such as, I am not sure who else
>> from Google is working
>> on cgroups stuff), so feel free to forward this to anyone you believe
>> should go,

>> or just let me know whom I missed.

>>

>> Regards,

> > Kir.
>
> I just came up with the following preliminary list of sessions:
>
> <http://wiki.linuxplumbersconf.org/2012:containers>

Thanks. We could also add devices namespace to the list. I'm fine NOT having it in the list, as I think we're agreed that it should come after user namespaces are complete.

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Wed, 25 Jul 2012 11:52:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 07/25/2012 03:51 PM, Serge Hallyn wrote:
> Quoting Glauber Costa (glommer@parallels.com):
>> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:
>>> Gentlemen,
>>>
>>> We are organizing containers mini-summit during next Linux Plumbers (San
>>> Diego, August 29-31).
>>> The idea is to gather and discuss everything relevant to namespaces,
>>> cgroups, resource management,
>>> checkpoint-restore and so on.
>>>
>>> We are trying to come up with a list of topics to discuss, so please
>>> reply with topic suggestions, and
>>> let me know if you are going to come.
>>>
>>> I probably forgot a few more people (such as, I am not sure who else
>>> from Google is working
>>> on cgroups stuff), so feel free to forward this to anyone you believe
>>> should go,
>>> or just let me know whom I missed.
>>>
>>> Regards,
>>> Kir.
>>
>> I just came up with the following preliminary list of sessions:
>>
>> <http://wiki.linuxplumbersconf.org/2012:containers>
>
> Thanks. We could also add devices namespace to the list. I'm fine NOT
> having it in the list, as I think we're agreed that it should come after
> user namespaces are complete.
>
> How about a "future of namespaces" generic placeholder?

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Serge E. Hallyn](#) on Wed, 25 Jul 2012 12:15:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Glauber Costa (glommer@parallels.com):

> On 07/25/2012 03:51 PM, Serge Hallyn wrote:

> > Quoting Glauber Costa (glommer@parallels.com):

> >> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:

> >>> Gentlemen,

> >>>

> >>> We are organizing containers mini-summit during next Linux Plumbers (San

> >>> Diego, August 29-31).

> >>> The idea is to gather and discuss everything relevant to namespaces,

> >>> cgroups, resource management,

> >>> checkpoint-restore and so on.

> >>>

> >>> We are trying to come up with a list of topics to discuss, so please

> >>> reply with topic suggestions, and

> >>> let me know if you are going to come.

> >>>

> >>> I probably forgot a few more people (such as, I am not sure who else

> >>> from Google is working

> >>> on cgroups stuff), so feel free to forward this to anyone you believe

> >>> should go,

> >>> or just let me know whom I missed.

> >>>

> >>> Regards,

> >>> Kir.

> >>

> >> I just came up with the following preliminary list of sessions:

> >>

> >> <http://wiki.linuxplumbersconf.org/2012:containers>

> >

> > Thanks. We could also add devices namespace to the list. I'm fine NOT

> > having it in the list, as I think we're agreed that it should come after

> > user namespaces are complete.

> >

> > How about a "future of namespaces" generic placeholder?

Sure. And I suppose syslog and time ns could also fall under that. Issues to discuss include:

1. what other namespaces do people expect to need?
2. justifications for the proposed namespaces, and other potential ways to address each. I.e., for devices ns, two justifications are (1) per-container loop devices (could be pre-allocated with containers aware of which /dev/loopN index they may use), and (2) filtering uevents only to appropriate containers to limit the uevent storm when containers do udevadm trigger --action=add (not sure, but we can

brainstorm other solutions).

3. semantics for the namespace - I.e syslog namespace, how do we decide whether/when to unshare it, is it owned by a user namespace or a pid namespace, etc.

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [ebiederm](#) on Thu, 26 Jul 2012 03:57:46 GMT

[View Forum Message](#) <> [Reply to Message](#)

Glauber Costa <glommer@parallels.com> writes:

> I just came up with the following preliminary list of sessions:

>

> <http://wiki.linuxplumbersconf.org/2012:containers>

>

> Since people mostly said what they wanted to talk about, but without
> extensive descriptions, I took the liberty of coming up with a small
> text for each in the blueprints. If you believe this is inaccurate, or
> would like to see it extended (although I personally don't see the point
> about going into very formal and deep details here), just let me know
> and I will edit it.

>

> This is all still subject to change.

Something that just came up recently and worth looking at if it hasn't already be resolved.

The network namespace, the user namespace, and the memory control group are not meshing well.

In particular we need some additional checks for an unprivileged user who can set `tcp_mem`. If you are the creator of a network namespace you should at least be able to set the values down. I don't know at all about increasing the amount of memory consumed by the tcp stack.

The non-nesting nature of memory control groups with respect to the network stack also seems very bizarre.

Another old issue is that unless I have missed something control groups are still broken for generic use in containers. Does anyone care? Are there any plans on fixing this issue?

Eric

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Thu, 26 Jul 2012 09:16:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 07/26/2012 07:57 AM, Eric W. Biederman wrote:

> Glauber Costa <glommer@parallels.com> writes:

>

>> I just came up with the following preliminary list of sessions:

>>

>> <http://wiki.linuxplumbersconf.org/2012:containers>

>>

>> Since people mostly said what they wanted to talk about, but without
>> extensive descriptions, I took the liberty of coming up with a small
>> text for each in the blueprints. If you believe this is inaccurate, or
>> would like to see it extended (although I personally don't see the point
>> about going into very formal and deep details here), just let me know
>> and I will edit it.

>>

>> This is all still subject to change.

>

> Something that just came up recently and worth looking at if it hasn't
> already be resolved.

>

> The network namespace, the user namespace, and the memory control group
> are not meshing well.

>

> In particular we need some additional checks for an unprivileged user
> who can set tcp_mem. If you are the creator of a network namespace you
> should at least be able to set the values down. I don't know at all
> about increasing the amount of memory consumed by the tcp stack.

This is between the user namespace and net namespace only, right ?

To be quite honest, I haven't looked thoroughly at UNS after your last
work. How do you yourself believe this should be?

>

> The non-nesting nature of memory control groups with respect to the
> network stack also seems very bizarre.

Correction:

The non-nesting nature of memory control groups is very bizarre. No need
for modifiers. It does support nesting, though. Just that it is
selectable, and not the default. But there is work in progress to change
that.

>

> Another old issue is that unless I have missed something control groups

> are still broken for generic use in containers. Does anyone care?
> Are there any plans on fixing this issue?
>
> Eric
>

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Thu, 26 Jul 2012 09:26:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

>>
>> Another old issue is that unless I have missed something control groups
>> are still broken for generic use in containers. Does anyone care?
>> Are there any plans on fixing this issue?
>>

What is "generic use in containers" ? I am using them alright, but not sure if this counts as generic or specific =)

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [ebiederm](#) on Thu, 26 Jul 2012 10:42:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

Glauber Costa <glommer@parallels.com> writes:

>>>
>>> Another old issue is that unless I have missed something control groups
>>> are still broken for generic use in containers. Does anyone care?
>>> Are there any plans on fixing this issue?
>>>
>
> What is "generic use in containers" ? I am using them alright, but not
> sure if this counts as generic or specific =)

The general container use case would be.

- Create a new mount namespace.
- Create fresh mounts of all of the control groups like I would do at boot, with no consideration to any other control group state.
- Start forking processes.

The expected semantics would be something like chroot for control groups, where all of the control groups that are created by fresh mounts are relative to whatever state the process of being in a control group that the process that mounted them was in.

Last I looked the closest you could come to that was bind mounts, and even with bind mounts you get into weird things where control groups are bound into hierarchies and you may be running a distribution that wants it's control groups bound into different hierarchies.

Last I looked this was just about a total disaster, and the only thing that allowed systemd to run in containers was the fact that systemd did not use controllers.

Eric

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [ebiederm](#) on Thu, 26 Jul 2012 10:57:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

Glauber Costa <glommer@parallels.com> writes:

> On 07/26/2012 07:57 AM, Eric W. Biederman wrote:

>> Glauber Costa <glommer@parallels.com> writes:

>>

>>> I just came up with the following preliminary list of sessions:

>>>

>>> <http://wiki.linuxplumbersconf.org/2012:containers>

>>>

>>> Since people mostly said what they wanted to talk about, but without
>>> extensive descriptions, I took the liberty of coming up with a small
>>> text for each in the blueprints. If you believe this is inaccurate, or
>>> would like to see it extended (although I personally don't see the point
>>> about going into very formal and deep details here), just let me know
>>> and I will edit it.

>>>

>>> This is all still subject to change.

>>

>> Something that just came up recently and worth looking at if it hasn't
>> already be resolved.

>>

>> The network namespace, the user namespace, and the memory control group
>> are not meshing well.

>>

>> In particular we need some additional checks for an unprivileged user
>> who can set tcp_mem. If you are the creator of a network namespace you
>> should at least be able to set the values down. I don't know at all
>> about increasing the amount of memory consumed by the tcp stack.

>

> This is between the user namespace and net namespace only, right ?

>

> To be quite honest, I haven't looked thoroughly at UNS after your last
> work. How do you yourself believe this should be?

I looked a little deeper and there are a few more places in the networking stack besides tcp_mem, that are setting memory limits that unprivileged users should not be able to touch.

What I would expect one of.

- A global limit accross network namespaces that the per netns limit can not allow you to escape.
- Something like rlimits where the limit can be reduced but not increased.
- capability checks that prevent anyone except the global root from changing the value (of course this has the problem that creating a fresh network namespace allows memory limit escaping).

The driving factor is that in the 3.7 time frame it will be possible to create user namespaces and network namespaces as unprivileged users. So we have to be careful that we setup the limits so that the global root on the host can set limits that are not trivially overridden.

>> The non-nesting nature of memory control groups with respect to the
>> network stack also seems very bizarre.

>

> Correction:

>

> The non-nesting nature of memory control groups is very bizarre. No need
> for modifiers. It does support nesting, though. Just that it is
> selectable, and not the default. But there is work in progress to change
> that.

Which leads to another bit of fun. It is possible to create containers in containers in containers, which has interesting implications for control groups in general, and especially interesting implications for control groups that don't nest.

Eric

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [Andrea Righi](#) on Thu, 26 Jul 2012 11:16:44 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Wed, Jul 25, 2012 at 02:00:41PM +0400, Glauber Costa wrote:

> On 07/25/2012 02:00 PM, Eric W. Biederman wrote:

> > Glauber Costa <glommer@parallels.com> writes:

> >

> >> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:

> >>> Gentlemen,
> >>>
> >>> We are organizing containers mini-summit during next Linux Plumbers (San
> >>> Diego, August 29-31).
> >>> The idea is to gather and discuss everything relevant to namespaces,
> >>> cgroups, resource management,
> >>> checkpoint-restore and so on.
> >>>
> >>> We are trying to come up with a list of topics to discuss, so please
> >>> reply with topic suggestions, and
> >>> let me know if you are going to come.
> >>>
> >>> I probably forgot a few more people (such as, I am not sure who else
> >>> from Google is working
> >>> on cgroups stuff), so feel free to forward this to anyone you believe
> >>> should go,
> >>> or just let me know whom I missed.
> >>>
> >>> Regards,
> >>> Kir.
> >>
> >> BTW, sorry for not replying before (vacations + post-vacations laziness)
> >>
> >> I would be interested in adding /proc virtualization to the discussion.
> >> By now it seems userspace would be the best place for that to happen, in
> >> a fuse overlay. I know Daniel has an initial implementation of that, and
> >> it would be good to have it as library that both OpenVZ and LXC (and
> >> whoever else wants) can use.
> >>
> >> Shouldn't take much time...
> >
> > What would you need proc virtualization for?
> >
>
> proc provides a lot of information that userspace tools rely upon.
> For instance, when running top, you can draw per-process figures from
> what we have now, but you can't make sense of percentages without
> aggregating container-wide information.
>
> When you read /proc/cpuinfo, as well, you would expect to see something
> that matches your container configuration.
>
> "free" is another example. The list goes on.

Another interesting feature IMHO would be the per-cgroup loadavg. A typical use case could be a monitoring system that wants to know which containers are more overloaded than others, instead of using a single system-wide measure in /proc/loadavg.

-Andrea

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [Serge E. Hallyn](#) on Thu, 26 Jul 2012 18:09:08 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Eric W. Biederman (ebiederm@xmission.com):

> Glauber Costa <glommer@parallels.com> writes:

>

> > I just came up with the following preliminary list of sessions:

> >

> > <http://wiki.linuxplumbersconf.org/2012:containers>

> >

> > Since people mostly said what they wanted to talk about, but without
> > extensive descriptions, I took the liberty of coming up with a small
> > text for each in the blueprints. If you believe this is inaccurate, or
> > would like to see it extended (although I personally don't see the point
> > about going into very formal and deep details here), just let me know
> > and I will edit it.

> >

> > This is all still subject to change.

>

> Something that just came up recently and worth looking at if it hasn't
> already be resolved.

>

> The network namespace, the user namespace, and the memory control group
> are not meshing well.

>

> In particular we need some additional checks for an unprivileged user
> who can set tcp_mem. If you are the creator of a network namespace you
> should at least be able to set the values down. I don't know at all
> about increasing the amount of memory consumed by the tcp stack.

>

> The non-nesting nature of memory control groups with respect to the
> network stack also seems very bizarre.

>

>

> Another old issue is that unless I have missed something control groups
> are still broken for generic use in containers. Does anyone care?
> Are there any plans on fixing this issue?

Can you elaborate? Which specific breakages are you thinking of?

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [Tejun Heo](#) on Thu, 26 Jul 2012 18:15:28 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello, Eric.

On Thu, Jul 26, 2012 at 03:42:50AM -0700, Eric W. Biederman wrote:

- > - Create a new mount namespace.
- > - Create fresh mounts of all of the control groups like I would do at
- > boot, with no consideration to any other control group state.
- > - Start forking processes.
- >
- > The expected semantics would be something like chroot for control
- > groups, where all of the control groups that are created by fresh mounts
- > are relative to whatever state the process of being in a control group
- > that the process that mounted them was in.

No, any attempt to build namespace support into cgroup core code will be nacked with strong prejudice. I still think it was a mistake to add that to sysfs. Thankfully, procfs is going the FUSE way and I hope in time we could convert sysfs to a similar mechanism and deprecate the in-kernel support.

So, no, no, no, no, no, no, no, no, no, no, no, no, no. :P

--
tejun

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers

Posted by [Serge E. Hallyn](#) on Thu, 26 Jul 2012 18:16:29 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Eric W. Biederman (ebiederm@xmission.com):

- > Glauber Costa <glommer@parallels.com> writes:
- >
- > >>>
- > >>> Another old issue is that unless I have missed something control groups
- > >>> are still broken for generic use in containers. Does anyone care?
- > >>> Are there any plans on fixing this issue?
- > >>>
- > >
- > > What is "generic use in containers" ? I am using them alright, but not
- > > sure if this counts as generic or specific =)
- >
- > The general container use case would be.
- >
- > - Create a new mount namespace.
- > - Create fresh mounts of all of the control groups like I would do at

> boot, with no consideration to any other control group state.
> - Start forking processes.
>
> The expected semantics would be something like chroot for control
> groups, where all of the control groups that are created by fresh mounts
> are relative to whatever state the process of being in a control group
> that the process that mounted them was in.
>
> Last I looked the closest you could come to that was bind mounts, and
> even with bind mounts you get into weird things where control groups are
> bound into hierarchies and you may be running a distribution that wants
> it's control groups bound into different hierarchies.
>
> Last I looked this was just about a total disaster, and the only thing
> that allowed systemd to run in containers was the fact that systemd did
> not use controllers.
>
> Eric

(Sorry, please disregard my last email :)

Yes, what we do now in ubuntu quantal is the bind mounts you mention,
and only optionally (using a startup hook).

Each container is brought up in say

/sys/fs/cgroup/devices/lxc/container1/container1.real, and that dir is
bind-mounted under /sys/fs/cgroup/devices in the guest. The guest
is not allowed to mount cgroup fs himself.

It's certainly not ideal (and in cases where cgroup allows you to
raise your own limits, worthless). The 'fake cgroup root' has been
mentioned before to address this. Definately worth discussing.

thanks,
-serge

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [ebiederm](#) on Thu, 26 Jul 2012 19:19:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

Tejun Heo <tj@kernel.org> writes:

> Hello, Eric.
>
> On Thu, Jul 26, 2012 at 03:42:50AM -0700, Eric W. Biederman wrote:
>> - Create a new mount namespace.
>> - Create fresh mounts of all of the control groups like I would do at
>> boot, with no consideration to any other control group state.

>> - Start forking processes.

>>

>> The expected semantics would be something like chroot for control groups, where all of the control groups that are created by fresh mounts are relative to whatever state the process of being in a control group that the process that mounted them was in.

>

> No, any attempt to build namespace support into cgroup core code will be nacked with strong prejudice.

The cgroup code was only merged with the understanding that this support was simple to add and it would be added. I am sorry that no one had the sense to follow up and make certain that promise was not fulfilled.

> I still think it was a mistake to add that to sysfs.

sysfs fundamentally can not represent all of the network devices in the hierarchy of objects that it chose.

sysfs does not have namespace hacks. Sysfs has hacks for the limitations of the hierarchy of devices that was chosen for a sysfs user space ABI.

> Thankfully, procfs is going the FUSE way.

No procfs is not going the FUSE way. Hacks for programs that misuse information in procfs is going the FUSE way.

The best example is there currently is not a good method for programs to figure out how parallel it is productive to be so the programs read /proc/cpuinfo and get the count of cpus. Control groups can limit you to fewer cpus but those programs have figured that out yet.

But ultimately fuse for procfs is about the rare case where people want to lie to applications, because it is easier to lie to applications than to disabuse the applications of their mistaken assumptions.

I have not seen a single suggest that any of the other procfs bits can go away.

> and I hope in time we could convert sysfs to a similar mechanism and > deprecate the in-kernel support.

I have nothing that even suggests there is a reasonable possibility of using fuse to deprecate any of the proc or sysfs support.

> So, no, no, no, no, no, no, no, no, no, no, no, no. :P

Bahahahahahahaha! :P

I sort of wish I had the energy to tackle this. As it is control groups hierarchies have very severe usability problems supporting one of their core use cases.

We should have our interfaces designed such that it is possible to run nested init's without hacks, and the only significant piece left on the hacks pile is control groups.

Control group hierarchies are a really strange piece of work whose design makes very little sense to me.

I think all I want from control groups is that a process that is bound into a control group hierarchy when it mounts that hierarchy will get not the normal control group root but instead the dentry of the directory for that process's place in the control group hierarchy.

What is that maybe 15-30 lines of code to look up the right dentry?

Eric

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [ebiederm](#) on Thu, 26 Jul 2012 19:38:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

Serge Hallyn <serge.hallyn@canonical.com> writes:

> (Sorry, please disregard my last email :)

>

> Yes, what we do now in ubuntu quantal is the bind mounts you mention,
> and only optionally (using a startup hook).

> Each container is brought up in say

> /sys/fs/cgroup/devices/lxc/container1/container1.real, and that dir is

> bind-mounted under /sys/fs/cgroup/devices in the guest. The guest

> is not allowed to mount cgroup fs himself.

>

> It's certainly not ideal (and in cases where cgroup allows you to
> raise your own limits, worthless). The 'fake cgroup root' has been
> mentioned before to address this. Definitely worth discussing.

It is going to be interesting to see how all of the unprivileged operations work when the user-namespaces start allowing unprivileged users to do things (3.7 timeframe I hope).

I can see it making things both easier and harder. I would hope not actually being root will make it easier to keep from raising your own limits.

Running some operations as non-root will catch other places off guard where people were definitely expecting nothing of the kind.

There are a couple of networking memory limits exposed through sysctl that I don't expect we want everyone changing, that I need to figure out how to separate out from the rest. A concept that hasn't existed before.

Eric

Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Tejun Heo](#) on Thu, 26 Jul 2012 19:44:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hey, Eric.

On Thu, Jul 26, 2012 at 12:19:12PM -0700, Eric W. Biederman wrote:

> > No, any attempt to build namespace support into cgroup core code will
> > be nacked with strong prejudice.

>

> The cgroup code was only merged with the understanding that this support
> was simple to add and it would be added. I am sorry that no one had
> the sense to follow up and make certain that promise was not fulfilled.

Good chunk of cgroup is messy and I'm likely to continue to break a lot of whatever promises that have been made. :)

> > Thankfully, procfs is going the FUSE way.

>

> No procfs is not going the FUSE way. Hacks for programs that misuse
> information in procfs is going the FUSE way.

All those were proposed to be solved by "teaching" kernel procfs how to present itself differently.

> The best example is there currently is not a good method for programs
> to figure out how parallel it is productive to be so the programs
> read /proc/cpuinfo and get the count of cpus. Control groups can
> limit you to fewer cpus but those programs have figured that out yet.

Yeah, and you can handle that too nicely with FUSE. More on this later.

> But ultimately fuse for procfs is about the rare case where people
> want to lie to applications, because it is easier to lie to applications
> than to disabuse the applications of their mistaken assumptions.

I don't think so. They are necessary parts of representing a properly scoped environment.

> I have not seen a single suggest that any of the other procs bits
> can go away.

I think I made that a couple times now. I definitely intend to push things that way. Or, at least, I'll bark as hard as I can against adding more namespace stuff to system pseudo filesystems.

> > and I hope in time we could convert sysfs to a similar mechanism and
> > deprecate the in-kernel support.

>

> I have nothing that even suggests there is a reasonable possibility of
> using fuse to deprecate any of the proc or sysfs support.

Why not? If there's some deficiency in FUSE or notification mechanisms in pseudo FSes, let's fix them.

> > So, no, no, no, no, no, no, no, no, no, no, no, no, no, no, no. :P

>

> Bahahahahahahaha! :P

:)

> I sort of wish I had the energy to tackle this. As it is control groups
> hierarchies have very severe usability problems supporting one of their
> core use cases.

>

> We should have our interfaces designed such that it is possible to run
> nested init's without hacks, and the only significant piece left on the
> hacks pile is control groups.

>

> Control group hierarchies are are really strange piece of work whose
> design makes very little sense to me.

cgroups is riddled with confused designs but this is not it. The confusion is that namespace should play a major role in the design of system pseudo filesystems and that it can be achieved by playing peekaboo with dentries.

It obfuscates the code for niche use case - which in itself could be acceptable if that's the only / best way to achieve that - while not even being able to serve the said use case properly.

Thanks.

--
tejun

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Daniel Wagner](#) on Thu, 02 Aug 2012 08:37:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

On 11.07.2012 23:41, Kir Kolyshkin wrote:

> Gentlemen,
>
> We are organizing containers mini-summit during next Linux Plumbers (San
> Diego, August 29-31).
> The idea is to gather and discuss everything relevant to namespaces,
> cgroups, resource management,
> checkpoint-restore and so on.
>
> We are trying to come up with a list of topics to discuss, so please
> reply with topic suggestions, and
> let me know if you are going to come.

I'd like to give a short presentation on what we would like to do with cgroups and ConnMan including a demonstration. The demo will show what will happen when orthogonal cgroup hierarchies are getting used by more than just systemd. Also I'd like to trigger some discussion on the future of the networking controllers (net_prio, net_cls).

cheers,
daniel

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Glauber Costa](#) on Thu, 02 Aug 2012 08:37:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 08/02/2012 12:37 PM, Daniel Wagner wrote:

> Hi,
>
> On 11.07.2012 23:41, Kir Kolyshkin wrote:
>> Gentlemen,
>>
>> We are organizing containers mini-summit during next Linux Plumbers (San
>> Diego, August 29-31).
>> The idea is to gather and discuss everything relevant to namespaces,
>> cgroups, resource management,

>> checkpoint-restore and so on.
>>
>> We are trying to come up with a list of topics to discuss, so please
>> reply with topic suggestions, and
>> let me know if you are going to come.
>
> I'd like to give a short presentation on what we would like to do with
> cgroups and ConnMan including a demonstration. The demo will show what
> will happen when orthogonal cgroup hierarchies are getting used by more
> than just systemd. Also I'd like to trigger some discussion on the
> future of the networking controllers (net_prio, net_cls).
>

Daniel, could you fill up the details of this proposal in the plumbers system? I believe kir sent out the details on how to do it somewhere down this thread...

Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Daniel Wagner](#) on Thu, 02 Aug 2012 09:42:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Glauber,

On 02.08.2012 10:37, Glauber Costa wrote:
> On 08/02/2012 12:37 PM, Daniel Wagner wrote:
>> Hi,
>>
>> On 11.07.2012 23:41, Kir Kolyshkin wrote:
>>> Gentlemen,
>>>
>>> We are organizing containers mini-summit during next Linux Plumbers (San
>>> Diego, August 29-31).
>>> The idea is to gather and discuss everything relevant to namespaces,
>>> cgroups, resource management,
>>> checkpoint-restore and so on.
>>>
>>> We are trying to come up with a list of topics to discuss, so please
>>> reply with topic suggestions, and
>>> let me know if you are going to come.
>>
>> I'd like to give a short presentation on what we would like to do with
>> cgroups and ConnMan including a demonstration. The demo will show what
>> will happen when orthogonal cgroup hierarchies are getting used by more
>> than just systemd. Also I'd like to trigger some discussion on the
>> future of the networking controllers (net_prio, net_cls).
>>
>

> Daniel, could you fill up the details of this proposal in the plumbers
> system? I believe kir sent out the details on how to do it somewhere
> down this thread...

Thanks, just added this one here:

<https://blueprints.launchpad.net/lpc/+spec/lpc2012-cont-netw+ork>

cheers,
daniel
