

---

Subject: OLS paper topics

Posted by [serue](#) on Wed, 23 Jan 2008 19:57:36 GMT

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

---

Hi,

here is a list of topics which I believe people are interested in writing papers on. I'm listing names of those who I think are interested in writing them. Sorry if I leave anyone off of a topic they're interested in. However it seems to me it would be best if we can agree on one person or two people to drive each topic, so everyone doesn't sit around expecting someone else to submit the abstract.

Am I missing any?

mini-summit:

I will submit for a 1-day mini-summit. Some interesting remaining topics for the mini-summit would include device namespaces, ttys and syslog, and lots of checkpoint/restart.

Does anyone think we don't need one of these at ols? Or that we do?

Is anyone interested in organizing the summit - coming out with an agenda, sending out announcements, etc - either alone or with my help?

pidns: (Pavel and Suka)

I've heard it called a tutorial, though I think some of the technical details are interesting in and of themselves. Its also an important area to make sure other developers - i.e people working with flocks or kthreads - understand.

netns: denis driving, daniel, benjamin

namespaces status: Pavel and Cedric

There was no ns status update last year it may be of interest. Instead of a separate pidns paper, pidns could also be mentioned here.

namespace entering: Cedric and serge?

This \*probably\* isn't enough for a full paper. So it could go under namespace status paper. But there is quite a bit to say just by listing the existing proposed solutions (at least 4 I can think of offhand) and their shortcomings.

memory c/r: Dave Hansen, serge interested

I suspect many people on this list have their own ideas on how to go about the checkpoint and restart. I suppose they

could each write their own paper, or work together on a single combined paper laying out the possibilities

user namespace approaches: serge

cgroups and containers: Paul Menage driving?, Balbir?  
A cgroups update could either be its own paper or joined with the namespaces status paper.

Paul were you considering a separate paper to discuss the cgroups and namespace management as laid out in your Sep 03 2007 email "Thoughts on Namespace / Subsystem unification"?

thanks,  
-serge

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [kir](#) on Thu, 24 Jan 2008 22:05:31 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Serge E. Hallyn wrote:

> Hi,  
>  
> here is a list of topics which I believe people are interested in  
> writing papers on. I'm listing names of those who I think are  
> interested in writing them. Sorry if I leave anyone off of a topic  
> they're interested in. However it seems to me it would be best if we  
> can agree on one person or two people to drive each topic, so everyone  
> doesn't sit around expecting someone else to submit the abstract.  
>  
> Am I missing any?  
>  
> mini-summit:  
> I will submit for a 1-day mini-summit. Some interesting  
> remaining topics for the mini-summit would include device  
> namespaces, ttys and syslog, and lots of checkpoint/restart.  
>

That's right, obviously the title of the mini-summit would be something like "Linux Kernel Containers".

> Does anyone think we don't need one of these at ols? Or that

- > we do?
- >
- > Is anyone interested in organizing the summit - coming out
- > with an agenda, sending out announcements, etc - either
- > alone or with my help?
- >

Guess I can help a bit with organizing this. To that effort, I have put up a wiki page:

[http://wiki.openvz.org/Containers/Mini-summit\\_2008](http://wiki.openvz.org/Containers/Mini-summit_2008)

We also need to have some kind of a list of attendees. So far I came with 12 names listed on that page, please feel free to edit/add more.

- > pidns: (Pavel and Suka)
- > I've heard it called a tutorial, though I think some of the
- > technical details are interesting in and of themselves. Its
- > also an important area to make sure other developers - i.e
- > people working with flocks or kthreads - understand.
- >

This is the proposal Pavel filed today, it is editable so we can improve it, please send your suggestion/fixes.

- > PID namespaces in the Linux kernel
- >
- > PID namespaces is a relatively new Linux kernel feature merged in
- > 2.6.24 kernel. It is a "view" of a particular set of tasks on the
- > system. PID namespaces work in a similar way to filesystem namespaces:
- > a process can be accessed in multiple namespaces, but it may have a
- > different name in each. It is one of the building blocks for
- > containers virtualization, and a prerequisite for
- > checkpointing/restart and live migration.
- >
- > The paper outlines some implementation details, explains user space
- > constraints that may seem odd, and discusses the impact of the feature
- > on the kernel APIs.
- >
- > In collaboration with Sukadev Bhattiprolu, IBM.

- > netns: denis driving, daniel, benjamin
- >

Right, Den Lunev, Daniel Lezcano, Pavel Emelyanov and Benjamin Thery. Den already filed a proposal for a paper/talk, here is how it looks like. Again, it is editable, so send your improvements.

- > Network namespace for Linux
- >

> The paper outlines the effort to implement a network virtualization in  
> the Linux kernel. This is a part of on-going effort to bring the  
> containers functionality into Linux. A container is an isolated  
> user-space partition, which performs like a stand-alone server, with  
> multiple containers co-existing on a single Linux box. Containers can  
> be used for resource management, network security and in  
> high-performance computing.

>  
> Making several instances of the Linux network stack, based on the  
> namespace concept, is a big challenge, but it is required to build a  
> full featured container. We will show how to configure and use a new  
> instance of the network stack, how the feature is architected and  
> implemented, and what is the current state of the art.

>  
> In collaboration with Daniel Lezcano, IBM, Benjamin Thery, Bull, and  
> Pavel Emelyanov, OpenVZ.

>>

> namespaces status: Pavel and Cedric  
> There was no ns status update last year it may be of  
> interest. Instead of a separate pidns paper, pidns could  
> also be mentioned here.

>

What if we organise a BoF, outlining the current status and future  
directions. Something like "Linux Kernel Containers development status"  
or some better title. I'd say "Containers" here instead of "Namespaces"  
(or use "Containers/Namespaces") because containers is easier term from  
my PoV.

> namespace entering: Cedric and serge?  
> This \*probably\* isn't enough for a full paper. So it could  
> go under namespace status paper. But there is quite a bit  
> to say just by listing the existing proposed solutions (at  
> least 4 I can think of offhand) and their shortcomings.

>

> memory c/r: Dave Hansen, serge interested  
> I suspect many people on this list have their own ideas on  
> how to go about the checkpoint and restart. I suppose they  
> could each write their own paper, or work together on a single  
> combined paper laying out the possibilities

>

Actually we already followed that way -- Andrey Mirkin has filed a  
paper/talk proposal today, titled "Containers checkpointing and live  
migration". Guess Dave (and/or Oren Laadan, and/or Cedric, maybe  
somebody else as well) could come with their own talks/papers as well.

Still can't make up my mind if we need a BoF on the subject or not.

- > user namespace approaches: serge
- >
- > cgroups and containers: Paul Menage driving?, Balbir?
- > A cgroups update could either be its own paper or joined
- > with the namespaces status paper.
- >
- > Paul were you considering a separate paper to discuss
- > the cgroups and namespace management as laid out in
- > your Sep 03 2007 email "Thoughts on Namespace / Subsystem
- > unification"?
- >

Not too much stuff about resource management, i.e. user memory controller, kernel memory controller, other per-namespace limits etc. Or is it all covered by cgroups? Or it's not what we are currently targeting?

Regards, Kir.

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [serue](#) on Thu, 24 Jan 2008 22:50:02 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Quoting Kir Kolyshkin (kir@openvz.org):

- >
- >
- > Serge E. Hallyn wrote:
- > > Hi,
- > >
- > > here is a list of topics which I believe people are interested in
- > > writing papers on. I'm listing names of those who I think are
- > > interested in writing them. Sorry if I leave anyone off of a topic
- > > they're interested in. However it seems to me it would be best if we
- > > can agree on one person or two people to drive each topic, so everyone
- > > doesn't sit around expecting someone else to submit the abstract.
- > >
- > > Am I missing any?
- > >
- > > mini-summit:
- > > I will submit for a 1-day mini-summit. Some interesting
- > > remaining topics for the mini-summit would include device
- > > namespaces, ttys and syslog, and lots of checkpoint/restart.

> >  
>  
> That's right, obviously the title of the mini-summit would be something  
> like "Linux Kernel Containers".

Attached is what I currently had for a proposal to send to the OLS  
committee. (Please feel free to edit and add to the wiki)

I just took guesses at moderator names because the OLS CFP asks for  
them. If someone else (Kir? Cedric?) wants to moderate the namespaces  
part that's fine with me.

> > Does anyone think we don't need one of these at ols? Or that  
> > we do?  
> >  
> > Is anyone interested in organizing the summit - coming out  
> > with an agenda, sending out announcements, etc - either  
> > alone or with my help?  
> >  
>  
> Guess I can help a bit with organizing this. To that effort, I have put  
> up a wiki page:  
> [http://wiki.openvz.org/Containers/Mini-summit\\_2008](http://wiki.openvz.org/Containers/Mini-summit_2008)

Cool, thanks.

> We also need to have some kind of a list of attendees. So far I came  
> with 12 names listed on that page, please feel free to edit/add more.  
> > pidns: (Pavel and Suka)  
> > I've heard it called a tutorial, though I think some of the  
> > technical details are interesting in and of themselves. Its  
> > also an important area to make sure other developers - i.e  
> > people working with flocks or kthreads - understand.  
> >  
> This is the proposal Pavel filed today, it is editable so we can improve  
> it, please send your suggestion/fixes.  
>  
> > PID namespaces in the Linux kernel  
> >  
> > PID namespaces is a relatively new Linux kernel feature merged in

Hmm... "PID namespaces are a relatively new Linux kernel feature"  
sounds more normal. Though I'm not sure which is more "correct"

> > 2.6.24 kernel. It is a "view" of a particular set of tasks on the  
> > system. PID namespaces work in a similar way to filesystem namespaces:  
> > a process can be accessed in multiple namespaces, but it may have a  
> > different name in each. It is one of the building blocks for

> > containers virtualization, and a prerequisite for  
> > checkpointing/restart and live migration.  
> >  
> > The paper outlines some implementation details, explains user space  
> > constraints that may seem odd, and discusses the impact of the feature  
> > on the kernel APIs.  
> >  
> > In collaboration with Sukadev Bhattiprolu, IBM.

Looks good to me.

>  
>  
> > netns: denis driving, daniel, benjamin  
> >  
>  
> Right, Den Lunev, Daniel Lezcano, Pavel Emelyanov and Benjamin Thery.  
> Den already filed a proposal for a paper/talk, here is how it looks  
> like. Again, it is editable, so send your improvements.  
>  
> > Network namespace for Linux  
> >  
> > The paper outlines the effort to implement a network virtualization in  
> > the Linux kernel. This is a part of on-going effort to bring the  
> > containers functionality into Linux. A container is an isolated  
> > user-space partition, which performs like a stand-alone server, with  
> > multiple containers co-existing on a single Linux box. Containers can  
> > be used for resource management, network security and in  
> > high-performance computing.  
> >  
> > Making several instances of the Linux network stack, based on the  
> > namespace concept, is a big challenge, but it is required to build a  
> > full featured container. We will show how to configure and use a new  
> > instance of the network stack, how the feature is architected and  
> > implemented, and what is the current state of the art.  
> >  
> > In collaboration with Daniel Lezcano, IBM, Benjamin Thery, Bull, and  
> > Pavel Emelyanov, OpenVZ.  
> >>  
>  
> > namespaces status: Pavel and Cedric  
> > There was no ns status update last year it may be of  
> > interest. Instead of a separate pidns paper, pidns could  
> > also be mentioned here.  
> >  
>  
> What if we organise a BoF, outlining the current status and future  
> directions. Something like "Linux Kernel Containers development status"

> or some better title. I'd say "Containers" here instead of "Namespaces"  
> (or use "Containers/Namespaces") because containers is easier term from  
> my PoV.

That has a different effect. A BoF would also be good, but if there are parts about the direction with namespaces about which we want some guidance from the community (which I think there are - sysfs, namespace entering, checkpoint/restart in general) then we may get more people at a paper talk than a bof. A Bof will basically get people particularly interested in either using or developing the feature.

>  
> > namespace entering: Cedric and serge?  
> > This \*probably\* isn't enough for a full paper. So it could  
> > go under namespace status paper. But there is quite a bit  
> > to say just by listing the existing proposed solutions (at  
> > least 4 I can think of offhand) and their shortcomings.  
> >  
> > memory c/r: Dave Hansen, serge interested  
> > I suspect many people on this list have their own ideas on  
> > how to go about the checkpoint and restart. I suppose they  
> > could each write their own paper, or work together on a single  
> > combined paper laying out the possibilities  
> >  
>  
> Actually we already followed that way -- Andrey Mirkin has filed a  
> paper/talk proposal today, titled "Containers checkpointing and live  
> migration". Guess Dave (and/or Oren Laadan, and/or Cedric, maybe  
> somebody else as well) could come with their own talks/papers as well.  
>  
> Still can't make up my mind if we need a BoF on the subject or not.

I figure at least a third of the mini-summit will be c/r. Separate papers may actually be the way to go, so long as each paper presents a different approach. OLS could put them all together in one block. Then at a BoF or a beer bof, after all have been presented and everyone has heard all the arguments, we can discuss the way to go forward.

>  
> > user namespace approaches: serge  
> >  
> > cgroups and containers: Paul Menage driving?, Balbir?  
> > A cgroups update could either be its own paper or joined  
> > with the namespaces status paper.  
> >  
> > Paul were you considering a separate paper to discuss  
> > the cgroups and namespace management as laid out in  
> > your Sep 03 2007 email "Thoughts on Namespace / Subsystem



> > unification"?

> >

>

> Not too much stuff about resource management, i.e. user memory

> controller, kernel memory controller, other per-namespace limits etc. Or

> is it all covered by cgroups? Or it's not what we are currently targeting?

I was figuring that each of those cgroups wouldn't have enough material for a paper and yes i figured one cgroup paper would be about the various cgroups. But I'm pretty far out of touch with that work so could be completely off base. The 'ns/cgroup unification'/administration topic is the one that interests me the most out of that block :)

thanks,

-serge

## Linux Kernel Containers Developer Summit

### \* Topic:

Namespaces, containers, cgroups, and container checkpoint restart.

### \* Description:

Development for namespaces and cgroups is well underway in the upstream kernel. To keep momentum going and keep the loosely knit teams of developers well-coordinated, a physical meeting in which to discuss future development plans is needed.

Additionally, we are at a point where crucial decisions about the nature of a "container object" and about the checkpoint restart design need to be made.

A final set of topics will be decided upon through mailing lists ahead of time, but potential topics include:

Handling filesystem/namespace synchronization

Handling of /proc and /sysfs within containers

Additional needed namespaces (i.e. Device namespace)

Nature of a 'container' - kernel object or userspace fiction

Additional cgroups and their design

How to initiate and synchronize checkpoint/restart

### \* Moderators:

(These are just suggestions, please let me know if you do or do not want to moderate)

Namespaces/containers: Serge Hallyn  
cgroups: Paul Menage  
Checkpoint/restart: Dave Hansen

\* Expected Time Required:

One day

\* Number of Attendees:

approx 30

\* Technical Requirements:

- \* Projector and screen
- \* whiteboard
- \* network access
- \* power for laptops.

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

## File Attachments

1) [minisummit.proposal](#), downloaded 394 times

---

---

Subject: Re: OLS paper topics  
Posted by [kir](#) on Fri, 25 Jan 2008 00:13:35 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Serge E. Hallyn wrote:

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

>

>> Serge E. Hallyn wrote:

>>

>>> Hi,

>>>

>>> here is a list of topics which I believe people are interested in  
>>> writing papers on. I'm listing names of those who I think are  
>>> interested in writing them. Sorry if I leave anyone off of a topic  
>>> they're interested in. However it seems to me it would be best if we  
>>> can agree on one person or two people to drive each topic, so everyone  
>>> doesn't sit around expecting someone else to submit the abstract.

>>>

>>> Am I missing any?

>>>

>>> mini-summit:  
>>> I will submit for a 1-day mini-summit. Some interesting  
>>> remaining topics for the mini-summit would include device  
>>> namespaces, ttys and syslog, and lots of checkpoint/restart.  
>>>  
>>>  
>> That's right, obviously the title of the mini-summit would be something  
>> like "Linux Kernel Containers".  
>>  
>  
> Attached is what I currently had for a proposal to send to the OLS  
> committee. (Please feel free to edit and add to the wiki)  
>

Looks really good from the first glance.

I put it to wiki as  
[http://wiki.openvz.org/Containers/Mini-summit\\_2008/Proposal](http://wiki.openvz.org/Containers/Mini-summit_2008/Proposal) and hope we  
will take a closer look tomorrow (as it's 3am here now:)).

> I just took guesses at moderator names because the OLS CFP asks for  
> them. If someone else (Kir? Cedric?) wants to moderate the namespaces  
> part that's fine with me.

>  
>

>>> Does anyone think we don't need one of these at ols? Or that  
>>> we do?

>>>  
>>> Is anyone interested in organizing the summit - coming out  
>>> with an agenda, sending out announcements, etc - either  
>>> alone or with my help?

>>>

>>>

>> Guess I can help a bit with organizing this. To that effort, I have put  
>> up a wiki page:

>> [http://wiki.openvz.org/Containers/Mini-summit\\_2008](http://wiki.openvz.org/Containers/Mini-summit_2008)

>>

>

> Cool, thanks.

>

>

>> We also need to have some kind of a list of attendees. So far I came  
>> with 12 names listed on that page, please feel free to edit/add more.

>>

>>> pidns: (Pavel and Suka)

>>> I've heard it called a tutorial, though I think some of the  
>>> technical details are interesting in and of themselves. Its  
>>> also an important area to make sure other developers - i.e

>>> people working with flocks or kthreads - understand.  
>>>  
>>>  
>> This is the proposal Pavel filed today, it is editable so we can improve  
>> it, please send your suggestion/fixes.  
>>  
>>  
>>> PID namespaces in the Linux kernel  
>>>  
>>> PID namespaces is a relatively new Linux kernel feature merged in  
>>>  
>  
> Hmm... "PID namespaces are a relatively new Linux kernel feature"  
> sounds more normal. Though I'm not sure which is more "correct"  
>  
Thanks! Fixed.  
>  
>>> 2.6.24 kernel. It is a "view" of a particular set of tasks on the  
>>> system. PID namespaces work in a similar way to filesystem namespaces:  
>>> a process can be accessed in multiple namespaces, but it may have a  
>>> different name in each. It is one of the building blocks for  
>>> containers virtualization, and a prerequisite for  
>>> checkpointing/restart and live migration.  
>>>  
>>> The paper outlines some implementation details, explains user space  
>>> constraints that may seem odd, and discusses the impact of the feature  
>>> on the kernel APIs.  
>>>  
>>> In collaboration with Sukadev Bhattiprolu, IBM.  
>>>  
>  
> Looks good to me.  
>

Great, thanks!

>  
>>  
>>> netns: denis driving, daniel, benjamin  
>>>  
>>>  
>> Right, Den Lunev, Daniel Lezcano, Pavel Emelyanov and Benjamin Thery.  
>> Den already filed a proposal for a paper/talk, here is how it looks  
>> like. Again, it is editable, so send your improvements.  
>>  
>>  
>>> Network namespace for Linux  
>>>

```

>>> The paper outlines the effort to implement a network virtualization in
>>> the Linux kernel. This is a part of on-going effort to bring the
>>> containers functionality into Linux. A container is an isolated
>>> user-space partition, which performs like a stand-alone server, with
>>> multiple containers co-existing on a single Linux box. Containers can
>>> be used for resource management, network security and in
>>> high-performance computing.
>>>
>>> Making several instances of the Linux network stack, based on the
>>> namespace concept, is a big challenge, but it is required to build a
>>> full featured container. We will show how to configure and use a new
>>> instance of the network stack, how the feature is architected and
>>> implemented, and what is the current state of the art.
>>>
>>> In collaboration with Daniel Lezcano, IBM, Benjamin Thery, Bull, and
>>> Pavel Emelyanov, OpenVZ.
>>>
>>>>
>>>>
>>> namespaces status: Pavel and Cedric
>>> There was no ns status update last year it may be of
>>> interest. Instead of a separate pidns paper, pidns could
>>> also be mentioned here.
>>>
>>>
>> What if we organise a BoF, outlining the current status and future
>> directions. Something like "Linux Kernel Containers development status"
>> or some better title. I'd say "Containers" here instead of "Namespaces"
>> (or use "Containers/Namespaces") because containers is easier term from
>> my PoV.
>>
>
> That has a different effect. A BoF would also be good, but if there are
> parts about the direction with namespaces about which we want some
> guidance from the community (which I think there are - sysfs, namespace
> entering, checkpoint/restart in general) then we may get more people at
> a paper talk than a bof. A Bof will basically get people particularly
> interested in either using or developing the feature.
>

```

Makes sense indeed.

My primary concern about talk vs. BoF was/is -- do we have enough quality material/content to write a "minimum of 6 pages and a maximum of 15 pages (properly formatted)" paper which is required for a talk (but not for a BoF). Well, a lot has happened since last year, maybe enough for a paper.

>  
>>> namespace entering: Cedric and serge?  
>>> This \*probably\* isn't enough for a full paper. So it could  
>>> go under namespace status paper. But there is quite a bit  
>>> to say just by listing the existing proposed solutions (at  
>>> least 4 I can think of offhand) and their shortcomings.  
>>>  
>>> memory c/r: Dave Hansen, serge interested  
>>> I suspect many people on this list have their own ideas on  
>>> how to go about the checkpoint and restart. I suppose they  
>>> could each write their own paper, or work together on a single  
>>> combined paper laying out the possibilities  
>>>  
>>>  
>> Actually we already followed that way -- Andrey Mirkin has filed a  
>> paper/talk proposal today, titled "Containers checkpointing and live  
>> migration". Guess Dave (and/or Oren Laadan, and/or Cedric, maybe  
>> somebody else as well) could come with their own talks/papers as well.  
>>  
>> Still can't make up my mind if we need a BoF on the subject or not.  
>>  
>  
> I figure at least a third of the mini-summit will be c/r. Separate  
> papers may actually be the way to go, so long as each paper presents a  
> different approach. OLS could put them all together in one block. Then  
> at a BoF or a beer bof, after all have been presented and everyone has  
> heard all the arguments, we can discuss the way to go forward.  
>

Sounds like a plan, especially the beer part ;)

>  
>>> user namespace approaches: serge  
>>>  
>>> cgroups and containers: Paul Menage driving?, Balbir?  
>>> A cgroups update could either be its own paper or joined  
>>> with the namespaces status paper.  
>>>  
>>> Paul were you considering a separate paper to discuss  
>>> the cgroups and namespace management as laid out in  
>>> your Sep 03 2007 email "Thoughts on Namespace / Subsystem  
>>> unification"?  
>>>  
>>>  
>> Not too much stuff about resource management, i.e. user memory  
>> controller, kernel memory controller, other per-namespace limits etc. Or  
>> is it all covered by cgroups? Or it's not what we are currently targeting?  
>>  
>

> I was figuring that each of those cgroups wouldn't have enough material  
> for a paper and yes i figured one cgroup paper would be about the  
> various cgroups. But I'm pretty far out of touch with that work so  
> could be completely off base. The 'ns/cgroup  
> unification'/administration topic is the one that interests me the most  
> out of that block :)  
>  
> thanks,  
> -serge  
>

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [kir](#) on Wed, 30 Jan 2008 08:50:59 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Serge E. Hallyn wrote:  
> Quoting Kir Kolyshkin (kir@openvz.org):  
>  
>> Serge E. Hallyn wrote:  
>>  
>>> mini-summit:  
>>> I will submit for a 1-day mini-summit. Some interesting  
>>> remaining topics for the mini-summit would include device  
>>> namespaces, ttys and syslog, and lots of checkpoint/restart.  
>>>  
>>>  
>> That's right, obviously the title of the mini-summit would be something  
>> like "Linux Kernel Containers".  
>>  
>  
> Attached is what I currently had for a proposal to send to the OLS  
> committee. (Please feel free to edit and add to the wiki)  
>

Some people have made a few minor changes in the version at wiki. I  
guess no more major changes will be done. Here it is:

== Topic ==

Namespaces, containers, cgroups, and container checkpoint/restart.

== Description ==

Development for namespaces and cgroups is well underway in the mainline kernel. To keep momentum going and keep the loosely knit teams of developers well-coordinated, a physical meeting in which to discuss future development plans is needed. Additionally, we are at a point where crucial decisions about the nature of a "container object" and about the checkpoint/restart design need to be made.

A final set of topics will be decided upon through mailing lists ahead of time, but potential topics include:

- \* Handling filesystem/namespaces synchronization
- \* Handling of /proc and /sysfs within containers
- \* Additional needed namespaces (i.e. device namespace)
- \* Nature of a 'container' — kernel object or userspace fiction
- \* Additional cgroups and their design
- \* How to initiate and synchronize checkpoint/restart
- \* How to enter a 'container' ?

== Moderators ==

{{Out|These are just suggestions, please edit here or let us know if you do or do not want to moderate.}}

- \* Namespaces/containers: Serge Hallyn, Cedric Le Goater
- \* cgroups: Paul Menage
- \* Checkpoint/restart: Dave Hansen, Cedric Le Goater

== Expected Time Required ==

One day.

== Number of Attendees ==

Approximately 30.

== Technical Requirements ==

- \* Projector and screen
- \* Whiteboard
- \* Network access
- \* Power for laptops
- \* doughnuts and croissants for participants

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---



Subject: Re: OLS paper topics  
Posted by [Balbir Singh](#) on Wed, 30 Jan 2008 08:58:13 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Kir Kolyshkin wrote:

> Serge E. Hallyn wrote:

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

>>

>>> Serge E. Hallyn wrote:

>>>

>>>> mini-summit:

>>>> I will submit for a 1-day mini-summit. Some interesting

>>>> remaining topics for the mini-summit would include device

>>>> namespaces, ttys and syslog, and lots of checkpoint/restart.

>>>>

>>> That's right, obviously the title of the mini-summit would be something

>>> like "Linux Kernel Containers".

>>>

>>

>> Attached is what I currently had for a proposal to send to the OLS

>> committee. (Please feel free to edit and add to the wiki)

>>

>

> Some people have made a few minor changes in the version at wiki. I

> guess no more major changes will be done. Here it is:

>

> == Topic ==

>

> Namespaces, containers, cgroups, and container checkpoint/restart.

>

> == Description ==

>

> Development for namespaces and cgroups is well underway in the

> mainline kernel. To keep momentum going and keep the loosely

> knit teams of developers well-coordinated, a physical meeting in

> which to discuss future development plans is needed.

> Additionally, we are at a point where crucial decisions about

> the nature of a "container object" and about the checkpoint/restart

> design need to be made.

>

> A final set of topics will be decided upon through mailing lists

> ahead of time, but potential topics include:

> \* Handling filesystem/namespace synchronization

> \* Handling of /proc and /sysfs within containers

> \* Additional needed namespaces (i.e. device namespace)

> \* Nature of a 'container' — kernel object or userspace fiction

> \* Additional cgroups and their design

> \* How to initiate and synchronize checkpoint/restart

> \* How to enter a 'container' ?

>

Resource Management for containers?

> == Moderators ==

> {{Out|These are just suggestions, please edit here or let us know if you  
> do or do not want to moderate.}}

>

> \* Namespaces/containers: Serge Hallyn, Cedric Le Goater

> \* cgroups: Paul Menage

> \* Checkpoint/restart: Dave Hansen, Cedric Le Goater

>

Resource Management? - I don't mind moderating that if it happens to be  
included in the summit or can we merge it with cgroups?

> == Expected Time Required ==

>

> One day.

>

> == Number of Attendees ==

>

> Approximately 30.

>

> == Technical Requirements ==

>

> \* Projector and screen

> \* Whiteboard

> \* Network access

> \* Power for laptops

> \* doughnuts and croissants for participants

>

---

> Containers mailing list

> Containers@lists.linux-foundation.org

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

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL

---

Containers mailing list

Containers@lists.linux-foundation.org

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

---

---

Subject: Re: OLS paper topics

Posted by [Cedric Le Goater](#) on Wed, 30 Jan 2008 09:53:26 GMT

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

---

>> == Topic ==

>>

>> Namespaces, containers, cgroups, and container checkpoint/restart.

>>

>> == Description ==

>>

>> Development for namespaces and cgroups is well underway in the  
>> mainline kernel. To keep momentum going and keep the loosely  
>> knit teams of developers well-coordinated, a physical meeting in  
>> which to discuss future development plans is needed.

>> Additionally, we are at a point where crucial decisions about  
>> the nature of a "container object" and about the checkpoint/restart  
>> design need to be made.

>>

>> A final set of topics will be decided upon through mailing lists  
>> ahead of time, but potential topics include:

>> \* Handling filesystem/namespace synchronization

>> \* Handling of /proc and /sysfs within containers

>> \* Additional needed namespaces (i.e. device namespace)

>> \* Nature of a 'container' — kernel object or userspace fiction

>> \* Additional cgroups and their design

>> \* How to initiate and synchronize checkpoint/restart

>> \* How to enter a 'container' ?

>>

>

> Resource Management for containers?

Yes we need that. There are a few possible conflicts in requirements that need to be discussed. For example, a resource management req would be to be able to move a process from one group to another and C/R wouldn't.

>> == Moderators ==

>> {{Out|These are just suggestions, please edit here or let us know if you  
>> do or do not want to moderate.}}

>>

>> \* Namespaces/containers: Serge Hallyn, Cedric Le Goater

>> \* cgroups: Paul Menage

>> \* Checkpoint/restart: Dave Hansen, Cedric Le Goater

IMO, 2 is a minimum to get the work done (it will require a minimum of organization) but the moderators don't necessarily have to be from the same team.

Or, having some one in charge of the logistics would be better ? In that

case, you can remove me from the moderators.

> Resource Management? - I don't mind moderating that if it happens to be  
> included in the summit or can we merge it with cgroups?

Good thing.

Cheers,

C.

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [Balbir Singh](#) on Wed, 30 Jan 2008 09:59:22 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Cedric Le Goater wrote:

>  
>>> == Topic ==  
>>>  
>>> Namespaces, containers, cgroups, and container checkpoint/restart.  
>>>  
>>> == Description ==  
>>>  
>>> Development for namespaces and cgroups is well underway in the  
>>> mainline kernel. To keep momentum going and keep the loosely  
>>> knit teams of developers well-coordinated, a physical meeting in  
>>> which to discuss future development plans is needed.  
>>> Additionally, we are at a point where crucial decisions about  
>>> the nature of a "container object" and about the checkpoint/restart  
>>> design need to be made.  
>>>  
>>> A final set of topics will be decided upon through mailing lists  
>>> ahead of time, but potential topics include:  
>>> \* Handling filesystem/namespace synchronization  
>>> \* Handling of /proc and /sysfs within containers  
>>> \* Additional needed namespaces (i.e. device namespace)  
>>> \* Nature of a 'container' — kernel object or userspace fiction  
>>> \* Additional cgroups and their design  
>>> \* How to initiate and synchronize checkpoint/restart  
>>> \* How to enter a 'container' ?  
>>>  
>>  
>> Resource Management for containers?

>  
> Yes we need that. There are a few possible conflicts in requirements  
> that need to be discussed. For example, a resource management req  
> would be to be able to move a process from one group to another and C/R  
> wouldn't.

Cedric,

Why is that a conflict w.r.t resource management, isn't that more of a cgroup feature - task migration. We need task migration for resource management, but task migration is a system administrator/management function. I don't see the conflict, but may be I am missing something very obvious.

>>> == Moderators ==  
>>> {{Out|These are just suggestions, please edit here or let us know if you  
>>> do or do not want to moderate.}}  
>>>  
>>> \* Namespaces/containers: Serge Hallyn, Cedric Le Goater  
>>> \* cgroups: Paul Menage  
>>> \* Checkpoint/restart: Dave Hansen, Cedric Le Goater  
>  
> IMO, 2 is a minimum to get the work done (it will require a minimum of  
> organization) but the moderators don't necessarily have to be from the  
> same team.  
>  
> Or, having some one in charge of the logistics would be better ? In that  
> case, you can remove me from the moderators.  
>  
>> Resource Management? - I don't mind moderating that if it happens to be  
>> included in the summit or can we merge it with cgroups?  
>  
> Good thing.  
>  
> Cheers,  
>  
> C.

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL

---

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

---

Subject: Re: OLS paper topics

Posted by [Dhaval Giani](#) on Wed, 30 Jan 2008 10:43:36 GMT

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

---

On Wed, Jan 30, 2008 at 02:28:13PM +0530, Balbir Singh wrote:

> Kir Kolyshkin wrote:

> > Serge E. Hallyn wrote:

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

> > >

> > > > Serge E. Hallyn wrote:

> > > >

> > > > mini-summit:

> > > > I will submit for a 1-day mini-summit. Some interesting

> > > > remaining topics for the mini-summit would include device

> > > > namespaces, ttys and syslog, and lots of checkpoint/restart.

> > > >

> > > That's right, obviously the title of the mini-summit would be something

> > > like "Linux Kernel Containers".

> > >

> > >

> > Attached is what I currently had for a proposal to send to the OLS

> > committee. (Please feel free to edit and add to the wiki)

> >

> >

> > Some people have made a few minor changes in the version at wiki. I

> > guess no more major changes will be done. Here it is:

> >

> > == Topic ==

> >

> > Namespaces, containers, cgroups, and container checkpoint/restart.

> >

> > == Description ==

> >

> > Development for namespaces and cgroups is well underway in the

> > mainline kernel. To keep momentum going and keep the loosely

> > knit teams of developers well-coordinated, a physical meeting in

> > which to discuss future development plans is needed.

> > Additionally, we are at a point where crucial decisions about

> > the nature of a "container object" and about the checkpoint/restart

> > design need to be made.

> >

> > A final set of topics will be decided upon through mailing lists

> > ahead of time, but potential topics include:

> > \* Handling filesystem/namespace synchronization

> > \* Handling of /proc and /sysfs within containers

> > \* Additional needed namespaces (i.e. device namespace)  
> > \* Nature of a 'container' — kernel object or userspace fiction  
> > \* Additional cgroups and their design  
> > \* How to initiate and synchronize checkpoint/restart  
> > \* How to enter a 'container' ?  
> >  
>  
> Resource Management for containers?  
>  
> > == Moderators ==  
> > {{Out|These are just suggestions, please edit here or let us know if you  
> > do or do not want to moderate.}}  
> >  
> > \* Namespaces/containers: Serge Hallyn, Cedric Le Goater  
> > \* cgroups: Paul Menage  
> > \* Checkpoint/restart: Dave Hansen, Cedric Le Goater  
> >  
>  
> Resource Management? - I don't mind moderating that if it happens to be  
> included in the summit or can we merge it with cgroups?  
>

I don't mind helping out here. Could you add me to the list as well?

Thanks,  
--  
regards,  
Dhaval

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [Cedric Le Goater](#) on Wed, 30 Jan 2008 12:57:20 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> Why is that a conflict w.r.t resource management, isn't that more of a cgroup  
> feature - task migration. We need task migration for resource management, but  
> task migration is a system administrator/management function. I don't see the  
> conflict, but may be I am missing something very obvious.

This is something we may want to dig. from a checkpoint/restart POV, a  
'container' acts a whole and you don't want tasks to be migrated from a  
container to another. They have to share the same container properties.

it might be possible to solve that issue from user space or by making sure

a set of cgroups are dedicated to a container.

C.

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [serue](#) on Wed, 30 Jan 2008 15:40:43 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Quoting Dhaval Giani (dhaval@linux.vnet.ibm.com):  
> On Wed, Jan 30, 2008 at 02:28:13PM +0530, Balbir Singh wrote:  
> > Kir Kolyshkin wrote:  
> > > Serge E. Hallyn wrote:  
> > > > Quoting Kir Kolyshkin (kir@openvz.org):  
> > > >  
> > > > Serge E. Hallyn wrote:  
> > > >  
> > > > mini-summit:  
> > > > I will submit for a 1-day mini-summit. Some interesting  
> > > > remaining topics for the mini-summit would include device  
> > > > namespaces, ttys and syslog, and lots of checkpoint/restart.  
> > > >  
> > > > That's right, obviously the title of the mini-summit would be something  
> > > > like "Linux Kernel Containers".  
> > > >  
> > > >  
> > > Attached is what I currently had for a proposal to send to the OLS  
> > > committee. (Please feel free to edit and add to the wiki)  
> > >  
> > >  
> > > Some people have made a few minor changes in the version at wiki. I  
> > > guess no more major changes will be done. Here it is:  
> > >  
> > > == Topic ==  
> > >  
> > > Namespaces, containers, cgroups, and container checkpoint/restart.  
> > >  
> > > == Description ==  
> > >  
> > > Development for namespaces and cgroups is well underway in the  
> > > mainline kernel. To keep momentum going and keep the loosely  
> > > knit teams of developers well-coordinated, a physical meeting in  
> > > which to discuss future development plans is needed.  
> > > Additionally, we are at a point where crucial decisions about



> > > the nature of a "container object" and about the checkpoint/restart  
> > > design need to be made.  
> > >  
> > > A final set of topics will be decided upon through mailing lists  
> > > ahead of time, but potential topics include:  
> > > \* Handling filesystem/namespace synchronization  
> > > \* Handling of /proc and /sysfs within containers  
> > > \* Additional needed namespaces (i.e. device namespace)  
> > > \* Nature of a 'container' ??? kernel object or userspace fiction  
> > > \* Additional cgroups and their design  
> > > \* How to initiate and synchronize checkpoint/restart  
> > > \* How to enter a 'container' ?  
> > >  
> >  
> >  
> > Resource Management for containers?  
> >  
> > > == Moderators ==  
> > > {{Out|These are just suggestions, please edit here or let us know if you  
> > > do or do not want to moderate.}}  
> > >  
> > > \* Namespaces/containers: Serge Hallyn, Cedric Le Goater  
> > > \* cgroups: Paul Menage  
> > > \* Checkpoint/restart: Dave Hansen, Cedric Le Goater  
> > >  
> >  
> > Resource Management? - I don't mind moderating that if it happens to be  
> > included in the summit or can we merge it with cgroups?  
> >  
>  
> I don't mind helping out here. Could you add me to the list as well?

Ok I figured resource mgmt and cgroups would be combined, but I'm fine  
splitting them. So can I do:

cgroups: Paul Menage, Balbir

Resource Management: dhaval, someone from openvz (Pavel?)

?

thanks,  
-serge

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

Subject: Re: OLS paper topics  
Posted by [serue](#) on Wed, 30 Jan 2008 15:48:01 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Quoting Cedric Le Goater (clg@fr.ibm.com):

>  
>> Why is that a conflict w.r.t resource management, isn't that more of a  
>> cgroup  
>> feature - task migration. We need task migration for resource management,  
>> but  
>> task migration is a system administrator/management function. I don't see  
>> the  
>> conflict, but may be I am missing something very obvious.  
>  
> This is something we may want to dig. from a checkpoint/restart POV, a  
> 'container' acts a whole and you don't want tasks to be migrated from a  
> container to another. They have to share the same container properties.  
> it might be possible to solve that issue from user space or by making sure  
> a set of cgroups are dedicated to a container.

This is certainly a good topic for the mini-summit.

The simple answer of course is you create a container for the virtual server, then create containers within those around the applications you want to checkpoint. But depending on logistics (i.e. per-container swapfiles) this could get complicated.

-serge

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [serue](#) on Wed, 30 Jan 2008 15:52:20 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Quoting Balbir Singh (balbir@linux.vnet.ibm.com):

> Cedric Le Goater wrote:  
> >  
> >>> == Topic ==  
> >>>  
> >>> Namespaces, containers, cgroups, and container checkpoint/restart.  
> >>>  
> >>> == Description ==  
> >>>  
> >>> Development for namespaces and cgroups is well underway in the

> >>> mainline kernel. To keep momentum going and keep the loosely  
> >>> knit teams of developers well-coordinated, a physical meeting in  
> >>> which to discuss future development plans is needed.  
> >>> Additionally, we are at a point where crucial decisions about  
> >>> the nature of a "container object" and about the checkpoint/restart  
> >>> design need to be made.  
> >>>  
> >>> A final set of topics will be decided upon through mailing lists  
> >>> ahead of time, but potential topics include:  
> >>> \* Handling filesystem/namespace synchronization  
> >>> \* Handling of /proc and /sysfs within containers  
> >>> \* Additional needed namespaces (i.e. device namespace)  
> >>> \* Nature of a 'container' ??? kernel object or userspace fiction  
> >>> \* Additional cgroups and their design  
> >>> \* How to initiate and synchronize checkpoint/restart  
> >>> \* How to enter a 'container' ?  
> >>>  
> >>  
> >> Resource Management for containers?  
> >  
> > Yes we need that. There are a few possible conflicts in requirements  
> > that need to be discussed. For example, a resource management req  
> > would be to be able to move a process from one group to another and C/R  
> > wouldn't.  
>  
> Cedric,  
>  
> Why is that a conflict w.r.t resource management, isn't that more of a cgroup  
> feature - task migration.

Oh, no - I think Cedric was saying that tasks need to be moved from  
container2 to container3 on the system for resourcemgmt - that is,  
reclassified, not migrated.

But for task migration we have resources tied to containers and  
processes are using those resources, so we can't migrate those tasks to  
another container in that case.

That's probably solved at the cgroup->container bindings, no? i.e.  
ns\_cgroup 2854 is my container and it's in a cpuset\_cgroup tied to cpus  
2 and 3, now i want to open it up to cpus 1-4, so i leave the tasks in  
ns\_cgroup 2854 but move it to another cpuset...

> We need task migration for resource management, but  
> task migration is a system administrator/management function. I don't see the  
> conflict, but may be I am missing something very obvious.  
>  
> >>> == Moderators ==

> >>> {{Out|These are just suggestions, please edit here or let us know if you  
> >>> do or do not want to moderate.}}  
> >>>  
> >>> \* Namespaces/containers: Serge Hallyn, Cedric Le Goater  
> >>> \* cgroups: Paul Menage  
> >>> \* Checkpoint/restart: Dave Hansen, Cedric Le Goater  
> >  
> > IMO, 2 is a minimum to get the work done (it will require a minimum of  
> > organization) but the moderators don't necessarily have to be from the  
> > same team.  
> >  
> > Or, having some one in charge of the logistics would be better ? In that  
> > case, you can remove me from the moderators.  
> >  
> >> Resource Management? - I don't mind moderating that if it happens to be  
> >> included in the summit or can we merge it with cgroups?  
> >  
> > Good thing.  
> >  
> > Cheers,  
> >  
> > C.  
>  
>  
> --  
> Warm Regards,  
> Balbir Singh  
> Linux Technology Center  
> IBM, ISTL  
>  
>  
> \_\_\_\_\_  
> Containers mailing list  
> Containers@lists.linux-foundation.org  
> <https://lists.linux-foundation.org/mailman/listinfo/containers>

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [Balbir Singh](#) on Wed, 30 Jan 2008 16:54:59 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Serge E. Hallyn wrote:  
> Quoting Balbir Singh (balbir@linux.vnet.ibm.com):  
>> Cedric Le Goater wrote:

```

>>>>> == Topic ==
>>>>>
>>>>> Namespaces, containers, cgroups, and container checkpoint/restart.
>>>>>
>>>>> == Description ==
>>>>>
>>>>> Development for namespaces and cgroups is well underway in the
>>>>> mainline kernel. To keep momentum going and keep the loosely
>>>>> knit teams of developers well-coordinated, a physical meeting in
>>>>> which to discuss future development plans is needed.
>>>>> Additionally, we are at a point where crucial decisions about
>>>>> the nature of a "container object" and about the checkpoint/restart
>>>>> design need to be made.
>>>>>
>>>>> A final set of topics will be decided upon through mailing lists
>>>>> ahead of time, but potential topics include:
>>>>> * Handling filesystem/namespace synchronization
>>>>> * Handling of /proc and /sysfs within containers
>>>>> * Additional needed namespaces (i.e. device namespace)
>>>>> * Nature of a 'container' ??? kernel object or userspace fiction
>>>>> * Additional cgroups and their design
>>>>> * How to initiate and synchronize checkpoint/restart
>>>>> * How to enter a 'container' ?
>>>>>
>>>> Resource Management for containers?
>>> Yes we need that. There are a few possible conflicts in requirements
>>> that need to be discussed. For example, a resource management req
>>> would be to be able to move a process from one group to another and C/R
>>> wouldn't.
>> Cedric,
>>
>> Why is that a conflict w.r.t resource management, isn't that more of a cgroup
>> feature - task migration.
>
> Oh, no - I think Cedric was saying that tasks need to be moved from
> container2 to container3 on the system for resourcemgmt - that is,
> reclassified, not migrated.
>

```

OK, so the concern is with automatic reclassification

```

> But for task migration we have resources tied to containers and
> processes are using those resources, so we can't migrate those tasks to
> another container in that case.
>
> That's probably solved at the cgroup->container bindings, no? i.e.
> ns_cgroup 2854 is my container and it's in a cpuset_cgroup tied to cpus
> 2 and 3, now i want to open it up to cpus 1-4, so i leave the tasks in

```

> ns\_cgroup 2854 but move it to another cpuset...  
>

Yes, I agree

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: OLS paper topics  
Posted by [Dhaval Giani](#) on Thu, 31 Jan 2008 07:49:16 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, Jan 30, 2008 at 09:40:43AM -0600, Serge E. Hallyn wrote:  
> Quoting Dhaval Giani (dhaval@linux.vnet.ibm.com):  
>  
> Ok I figured resource mgmt and cgroups would be combined, but I'm fine  
> splitting them. So can I do:  
>  
> cgroups: Paul Menage, Balbir  
> Resource Management: dhaval, someone from openvz (Pavel?)  
>  
> ?

Fine with me.

>  
> thanks,  
> -serge

--

regards,  
Dhaval

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---