

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

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

